arXiv:1603.01308v1 [stat.ME] 3 Mar 2016

Dynamic Adaptive Mixture Models

Leopoldo Catania[a,*]

[a]*Department of Economics and Finance, University of Rome, Tor Vergata, Rome, Italy*

Abstract

In this paper we propose a new class of Dynamic Mixture Models (DAMMs) being able to sequentially adapt the mixture components as well as the mixture composition using information coming from the data. The information driven nature of the proposed class of models allows to exactly compute the full likelihood and to avoid computer intensive simulation schemes. An extensive Monte Carlo experiment reveals that the new proposed model can accurately approximate the more complicated Stochastic Dynamic Mixture Model previously introduced in the literature as well as other kind of models. The properties of the new proposed class of models are discussed through the paper and an application in financial econometrics is reported.

Keywords: Dynamic mixture models, Score Driven models, adaptive models, density prediction.

1. Introduction

Mixtures of distributions are extremely diffused parametric tools used to model non–Gaussian shapes that usually characterise empirical data. A great level of flexibility can be achieved in Mixture Models (MM) by appropriate choices of the mixture components distributions. Moreover, the mixture components can also be adaptive with respect to the new information as usually happens in linear and generalised Mixture Models, see e.g. Bishop (2006). Within the context of MM, also the mixture composition can be allowed to evolve over time, this class of models is usually identified

*Department of Economics and Finance, University of Rome "Tor Vergata" , Via Columbia, 2, 00133, Rome, Italy. leopoldo.catania@uniroma2.it, Phone: +39 06 7259 5941, web page: http://www.economia.uniroma2.it/phd/ef/default.asp?a=216.

as Dynamic Mixture Models (DMM), see e.g. Yu (2012). DMM have been successfully applied in process monitoring (Yu, 2012), intervention detections (Gerlach et al., 2000), insurance losses (Frigessi et al., 2002) and graphical engineering (KaewTraKulPong and Bowden, 2002; Xie et al., 2005). A drawback of DMM is that, when nonlinear non–Gaussian specifications are assumed for the mixture components and for the evolution of the mixture composition, classical inference cannot be applied anymore, see e.g. Gerlach et al. (2000). Usual solutions relay on computer intensive Markov Chain Monte Carlo (MCMC) simulation schemes to carry out Bayesian inference which highly reduce the attractiveness of such models and their implementation in commercial softwares, see e.g. Gerlach et al. (2000), Yu (2012) and Billio et al. (2013).

In this paper we follow a different approach to model the time evolution of the mixture component distributions as well as the mixture composition in a fully observation driven framework (Cox et al., 1981). We build our model starting from recent advances in Score Driven models, see e.g. Harvey (2013) and Creal et al. (2013). In Score Driven models, the latent dynamic parameters are updated using an adequate forcing variable based on the score of the conditional distribution. In our context, the mixture components can be chosen to be any parametric distribution with the possibility of allowing for time variation in the full set of parameters of each component. We also allow for the mixture composition to be sequentially updated using the information contained in data. We call this class of models Dynamic Adaptive Mixture Models (DAMMs) given their high flexibility in terms of possible dynamic parametric assumptions and their ability to sequentially adapt the mixture composition. As extensively discussed by Koopman et al. (2015), the use of the conditional score in order to pool the information coming from the data into new updated parameters, results to be extremely flexible. Indeed, they found that robust score–based filters well approximate unobserved dynamics generated from nonlinear non–Gaussian state space models and other kinds of models.
Usually, in applications that exploit mixture of distributions as parametric tools, one of the main inferential objective is in terms of clustering results, see e.g. McLachlan and Peel (2000). Also within the context of Markov Switching (MS) models, empirical results are often given in terms of decoding for the latent unobservable regimes, see e.g. Hamilton (1989) , Cappé et al. (2005) and Frühwirth-

Schnatter (2006). However, given the dynamic evolution of each component of the mixture, the use of DAMM is not limited to this type of applications. Indeed, one of the main scope of DAMM is to adaptively represent the dynamic composition of the mixture, and not to uniquely identify regimes or structural changes affecting the data. However, if the underlying data generating process behaves as usual MS processes, then DAMM can easily accommodate this feature. In order to demonstrate the flexibility of DAMM we perform an extensive Monte Carlo experiment composed by four parts. In the first part, we aim to approximate the first two conditional moments as well as the dynamic mixture composition generated by a Stochastic DMM (SDMM) similar to that of Yu (2012). In this respect, our experiment is similar to that performed by Koopman et al. (2015) in the context of filtering state space models using Generalised Autoregressive Score (GAS) models, a theoretical treatment of this interesting topic is reported by Harvey (2013). The second and the third parts, similarly to Engle (2002) and Creal et al. (2013), focus on filtering several artificial patterns assumed for the correlation and the mixture composition of conditional random variables. In the last experiment we investigate the cost, in terms of precision loss in filtering conditional correlations, of model misspecification. Our results suggest that DAMMs are able to adequately approximate the highly nonlinear dynamics generated by the SDMM and the artificial correlation and mixture composition patterns better then competitive models. To further investigate the properties of the proposed DAMMs, we also report an empirical application in financial econometrics. Specifically, we estimate several univariate and multivariate DAMM specifications to a panel of financial returns. The specifications we consider differ in terms of parametric assumptions and dynamic properties of the conditional distribution. In sample and out of sample comparative results are given in terms of goodness of fit and predictive ability of the marginal and joint conditional distributions. We found that DAMMs outperforms competitive GARCH/DCC models in both cases.

The paper is organised in the following manner. Section 2 describes the DAMM and details the updating mechanism for the mixture component distributions and the mixture composition. Section 3 reports several univariate and multivariate DAMM specifications that can be used in relevant empirical applications. Section 4 reports the Monte Carlo experiments.

Sections 5 reports the empirical application in financial econometrics. Finally, Section 6 concludes and reports some suggestions for future research.

2. Dynamic Adaptive Mixture Models

Let $\mathbf{y}_t \in \Re^d$ be a d–dimensional random vector conditionally distributed according to $p(\mathbf{y}_t|\mathcal{F}_{t-1}, \boldsymbol{\theta}_t)$, with $\mathcal{F}_{t-1}$ be the filtration generated by the process $\{\mathbf{y}_s, s > 0\}$ up to time $t-1$, and $\boldsymbol{\theta}_t$ be a vector of time varying conditional parameters. We will assume $p(\cdot)$ to be a finite mixture of J real valued conditional distributions, i.e.

$$p(\mathbf{y}_t|\mathcal{F}_{t-1}, \boldsymbol{\theta}_t) = \sum_{j=1}^{J} \omega_{j,t} p_j(\mathbf{y}_t|\mathcal{F}_{t-1}, \boldsymbol{\theta}_{j,t}), \tag{1}$$

with $\omega_{j,t} \in (0,1)$ and $\sum_{j=1}^{J} \omega_{j,t} = 1 \quad \forall \quad t = 1, \ldots$ and $\boldsymbol{\theta}_t = \left(\boldsymbol{\theta}'_{j,t}, \omega_{j,t}, j = 1, \ldots, J\right)'$. Within the class of Dynamic Mixture Models, the mixture component density parameters $\boldsymbol{\theta}_{j,t}$, generally follow a stochastic process. Convenient choices are first order nonlinear autoregression (Billio et al., 2012; Casarin et al., 2015) and Markov Switching processes (Kim, 1994; Kim and Nelson, 1999; Ardia, 2008; Harrison and West, 1999). The latter are usually estimated by particle filters in a Bayesian context, while for the former the Expectation–Maximisation algorithm of Dempster et al. (1977) is frequently employed. Differently, in this paper we follow the Score Driven Framework (SDF) of Harvey (2013) and Creal et al. (2013) by letting the full set of parameters to be updated using the score of the conditional distribution $p(\mathbf{y}_t|\mathcal{F}_{t-1}, \boldsymbol{\theta}_t)$.

Formally, let $\boldsymbol{\omega}_t = (\omega_{j,t}, j = 1, \ldots, J)'$ be the vector containing the mixture weights at time t, and $\tilde{\boldsymbol{\omega}}_t \in \Re^{J-1}$ be a $(J-1)$–dimension vector such that $\Lambda^{\omega}(\tilde{\boldsymbol{\omega}}_t) = \boldsymbol{\omega}_t$, for a $\mathcal{F}_{t-1}$ measurable mapping function $\Lambda^{\omega} : \Re^{J-1} \to \mathcal{S}^J$, such that $\Lambda^{\omega} \in \mathcal{C}^2$, with $\mathcal{S}^J$ representing the standard unit J–simplex, i.e. $\mathcal{S}^J : \{(t_1, \ldots, t_J) \in \Re^J | \sum_{j=1}^{J} t_j = 1 \wedge t_j \geq 0, \quad \forall j\}$. Similarly, let $\tilde{\boldsymbol{\theta}}_{j,t} \in \Omega_j \subseteq \Re^{d_j}$ be a d_j–dimension vector such that, for each time t, we have $\Lambda^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right) = \boldsymbol{\theta}_{j,t}$ where $\Lambda^j : \Re^{d_j} \to \Omega_j$ holds the same properties stated for $\Lambda^{\omega}(\cdot)$, for all $j = 1, \ldots, J$. In order to avoid complicated nonlinear constraints on the parameters dynamic, in this paper, instead of directly modeling the vector $\boldsymbol{\theta}_t$ defined on

$\mathcal{S}^J \times \Omega^1 \times \cdots \times \Omega^J$, we model the unconstraint vector of parameters $\tilde{\boldsymbol{\theta}}_t = \left(\tilde{\boldsymbol{\omega}}_t', \tilde{\boldsymbol{\theta}}_j', j = 1, \ldots, J\right)'$ defined on $\Re^{J-1} \times \Re^{d_1} \times \cdots \times \Re^{d_J}$. To this end, we reparametrise the conditional distribution (1) into $\tilde{p}\left(\mathbf{y}_t|, \tilde{\boldsymbol{\theta}}_t\right)$, where henceforth the dependence from $\mathcal{F}_{t-1}$ has been omitted for notational purposes. In the SDF, the quantity of interest is the score of the conditional distribution given by

$$\tilde{\nabla}\left(\tilde{\boldsymbol{\theta}}_t|\mathbf{y}_t\right) = \left.\frac{\partial \ln \tilde{p}\left(\mathbf{y}_t|\tilde{\boldsymbol{\theta}}_t\right)}{\partial \tilde{\boldsymbol{\theta}}}\right|_{\tilde{\boldsymbol{\theta}}=\tilde{\boldsymbol{\theta}}_t}, \tag{2}$$

which enters linearly as a forcing variable into the dynamic updating equation of $\tilde{\boldsymbol{\theta}}_t$, i.e.

$$\tilde{\boldsymbol{\theta}}_{t+1} = \boldsymbol{\kappa} + \mathbf{A}\tilde{\nabla}\left(\tilde{\boldsymbol{\theta}}_t|\mathbf{y}_t\right) + \mathbf{B}\tilde{\boldsymbol{\theta}}_t, \tag{3}$$

where $\boldsymbol{\kappa}$ is a $\left(J - 1 + \sum_{j=1}^{J} d_j\right) = L$–dimension vector and $\mathbf{A}$ and $\mathbf{B}$ are $L \times L$ matrices of coefficients that need to be estimated. Since we are modeling the unconstraint vector of parameters $\tilde{\boldsymbol{\theta}}_t$ and $\{\tilde{\nabla}\left(\tilde{\boldsymbol{\theta}}_t|\mathbf{y}_s\right), s > 0\}$ forms a Martingale Difference Sequence, we only need to impose that the eigenvalues of $\mathbf{B}$ lies inside the unit circle to ensure weak stationarity of the process $\{\tilde{\boldsymbol{\theta}}_s, s > 0\}$. Constraints on the number of free parameters that are present in $\boldsymbol{\kappa}, \mathbf{A}$ and $\mathbf{B}$ can be imposed in order to avoid problems of parameters proliferation, indeed, throughout the paper we will impose a diagonal structure for $\mathbf{A}$ and $\mathbf{B}$. It is worth noting that, by a simple application of the chain rule, the conditional score with respect to the unconstraint vector of parameters $\tilde{\nabla}\left(\tilde{\boldsymbol{\theta}}_t|\mathbf{y}_t\right)$ can be easily represented as the product between the transpose of the jacobian of the full mapping function $\mathcal{J}\left(\tilde{\boldsymbol{\theta}}_t\right)$ and the conditional score evaluated with respect to the constraint vector of parameters $\nabla\left(\boldsymbol{\theta}_t|\mathbf{y}_t\right)$, i.e.,

$$\tilde{\nabla}\left(\tilde{\boldsymbol{\theta}}_t|\mathbf{y}_t\right) = \mathcal{J}\left(\tilde{\boldsymbol{\theta}}_t\right)' \nabla\left(\boldsymbol{\theta}_t|\mathbf{y}_t\right), \tag{4}$$

where with "full mapping function" we mean the vector–valued function $\Lambda : \Re^{J-1} \times \Re^{d_1} \times \cdots \times \Re^{d_J} \to \mathcal{S}^J \times \Omega^1 \times \cdots \times \Omega^J$ that incorporates $\Lambda^{\omega}(\cdot)$ and $\Lambda^{j}(\cdot), \quad j = 1, \ldots, J$, such that $\Lambda\left(\tilde{\boldsymbol{\theta}}_t\right) = \boldsymbol{\theta}_t, \quad \forall t$. In our context, since the matrices $\mathbf{A}$ and $\mathbf{B}$ are diagonal, and the matrix $\mathcal{J}\left(\tilde{\boldsymbol{\theta}}_t\right)$ is block diagonal, the dynamic updating equation (3) can be divided into $J + 1$ individual dynamics, i.e.

$$\tilde{\boldsymbol{\omega}}_{t+1} = \boldsymbol{\kappa}^{\omega} + \mathbf{A}^{\omega}\mathcal{J}^{\omega}\left(\tilde{\boldsymbol{\omega}}_t\right)' \nabla^{\omega}\left(\boldsymbol{\omega}_t|\mathbf{y}_t\right) + \mathbf{B}^{\omega}\tilde{\boldsymbol{\omega}}_t \tag{5}$$

$$\tilde{\boldsymbol{\theta}}_{j,t+1} = \boldsymbol{\kappa}^{j} + \mathbf{A}^{j}\mathcal{J}^{j}\left(\tilde{\boldsymbol{\theta}}_{j,t}\right)' \nabla^{j}\left(\boldsymbol{\theta}_{j,t}|\mathbf{y}_t\right) + \mathbf{B}^{j}\tilde{\boldsymbol{\theta}}_{j,t}, \quad j = 1, \ldots, J, \tag{6}$$

where all the symbols have the same interpretation as before, but are now related to each specific quantity of the model. The following Proposition will be necessary later.

Proposition 2.1 (**Score of a mixture of distributions**). *Let* $\mathbf{x} \sim g(\mathbf{x}|\boldsymbol{\phi})$, *with* $g(\mathbf{x}|\boldsymbol{\phi}) = \sum_{i=1}^{K} \omega_i g_i(\mathbf{x}|\boldsymbol{\phi}_i)$, *with* $\omega_i \geq 0$, $\sum_{i=1}^{K} \omega_i = 1$ *and* $\boldsymbol{\omega} = (\omega_i, i = 1, \ldots, K)'$. *Let* d_i *be the dimension of* $\boldsymbol{\phi}_i$ *and let* $\boldsymbol{\phi} = \left(\boldsymbol{\omega}', \boldsymbol{\phi}_i', i = 1, \ldots, K\right)'$ *be a* $\left(K + \sum_{i=1}^{K} d_i\right) = D$*–th dimension vector containing all the distribution parameters. Let*

$$\boldsymbol{\nabla}(\boldsymbol{\phi}|\mathbf{x}) = \frac{\partial \ln g(\mathbf{x}|\boldsymbol{\phi})}{\partial \boldsymbol{\phi}}, \tag{7}$$

be a vector of dimension D *representing the score of* $g(\mathbf{x}|\boldsymbol{\phi})$, *partitioned into* $\boldsymbol{\nabla}(\boldsymbol{\phi}|\mathbf{x}) = \left(\boldsymbol{\nabla}^{\boldsymbol{\omega}}(\boldsymbol{\omega}|\mathbf{x})', \boldsymbol{\nabla}^{(i)}(\boldsymbol{\phi}_i|\mathbf{x})', i = 1, \ldots, K\right)'$, *where*

$$\boldsymbol{\nabla}^{\boldsymbol{\omega}}(\boldsymbol{\omega}|\mathbf{x}) = \frac{\partial \ln g(\mathbf{x}|\boldsymbol{\phi})}{\partial \boldsymbol{\omega}} \tag{8}$$

$$\boldsymbol{\nabla}^{(i)}(\boldsymbol{\phi}_i|\mathbf{x}) = \frac{\partial \ln g(\mathbf{x}|\boldsymbol{\phi})}{\partial \boldsymbol{\phi}_i}, \quad i = 1, \ldots, K. \tag{9}$$

Assuming $g_i(\mathbf{x}|\boldsymbol{\phi}_i) \perp\!\!\!\perp g_l(\mathbf{x}|\boldsymbol{\phi}_l), \quad \forall \quad i \neq l$, *we have that*

(a) $\boldsymbol{\nabla}^{\boldsymbol{\omega}}(\boldsymbol{\omega}|\mathbf{x}) = \left(\frac{g_i(\mathbf{x}|\boldsymbol{\phi}_i)}{g(\mathbf{x}|\boldsymbol{\phi})}, i = 1, \ldots, K\right)'$, *and*

(b) $\boldsymbol{\nabla}^{(i)}(\boldsymbol{\phi}_i|\mathbf{x}) = \omega_i \frac{g_i(\mathbf{x}|\boldsymbol{\phi}_i)}{g(\mathbf{x}|\boldsymbol{\phi})} \boldsymbol{\nabla}_{g_i}(\boldsymbol{\phi}_i|\mathbf{x})$, $i = 1, \ldots, K$

where $\boldsymbol{\nabla}_{g_i}(\boldsymbol{\phi}_i|\mathbf{x}) = \frac{\partial \ln g_i(\mathbf{x}|\boldsymbol{\phi}_i)}{\partial \boldsymbol{\phi}_i}$ *is the score of the* i*–th mixture component distribution.*

Proof. (a) follows immediately, while for (b) we note that

$$\boldsymbol{\nabla}^{(i)}(\boldsymbol{\phi}_i|\mathbf{x}) = \omega_i \frac{1}{g(\mathbf{x}|\boldsymbol{\phi})} \frac{\partial g_i(\mathbf{x}|\boldsymbol{\phi}_i)}{\partial \boldsymbol{\phi}_i}, \tag{10}$$

and

$$\frac{\partial g_i(\mathbf{x}|\boldsymbol{\phi}_i)}{\partial \boldsymbol{\phi}_i} = \frac{\partial \exp(\ln g_i(\mathbf{x}|\boldsymbol{\phi}_i))}{\partial \boldsymbol{\phi}_i} = g_i(\mathbf{x}|\boldsymbol{\phi}_i) \boldsymbol{\nabla}_{g_i}(\boldsymbol{\phi}_i|\mathbf{x}). \tag{11}$$

□

2.1. Update the mixture composition

Several different choices are available in order to reparametrise and update the mixture weights $\boldsymbol{\omega}_t$. For example, Billio et al. (2013) use the Logistic–Transformed Gaussian (LTG) weights, i.e. their mapping function is the vector valued function $\Lambda^{\text{LTG}} : \Re^J \to \mathcal{S}^J$, with j–th component given by $\Lambda_j^{\text{LTG}}(\tilde{\omega}_{j,t}) = \frac{\exp(\tilde{\omega}_{j,t})}{\sum_{i=1}^J \exp(\tilde{\omega}_{i,t})}$, and $\tilde{\boldsymbol{\omega}}_t \sim \mathcal{N}_J(\tilde{\boldsymbol{\omega}}_t | \tilde{\boldsymbol{\omega}}_{t-1}, \Sigma)$, with Σ be a proper covariance matrix. This mapping and updating scheme assumes that the weights do not depend on the observable data. Moreover, it results to be somehow overparametrised since it is always possible to find a mapping function defined on $\Re^{J-1}$ (instead of $\Re^J$) with $J-1$ free parameters, which maps in $\mathcal{S}^J$. We propose a convenient choice for the function $\Lambda^{\omega}(\cdot)$, given by

$$\Lambda^{\omega}(\tilde{\boldsymbol{\omega}}_t) := \begin{cases} \omega_{j,t} = \lambda_{[0,b_{j,t}]}(\tilde{\omega}_{j,t}), & j = 1, \ldots, J-1 \\ \omega_{J,t} = 1 - \sum_{h=1}^{J-1} \omega_{h,t}, \end{cases} \tag{12}$$

where $b_{j,t} = b_{j-1,t} - \omega_{j-1,t}$ with $b_{1,t} = 1$ and $\lambda_{[L,U]} : \Re \to [L, U] \subset \Re$ is a real–valued deterministic monotone twice differentiable mapping function such as the modified logistic $\lambda_{[L,U]}(x) = L + \frac{(U-L)}{1+\exp(-x)}$. With these choices of $\Lambda^{\omega}(\cdot)$ and $\lambda_{[L,U]}(\cdot)$, the (j,h)–th element of the $J \times J-1$ jacobian matrix $\mathcal{J}^{\omega}(\cdot)$ is given by

$$\mathcal{J}^{\omega}(\tilde{\boldsymbol{\omega}}_t)_{(j,h)} = \begin{cases} \frac{b_{j,t}\exp(-\tilde{\omega}_{j,t})}{(1+\exp(-\tilde{\omega}_{j,t}))^2}, & \text{if } h = j \\ \frac{-\sum_{k=1}^{j-1} \mathcal{J}^{\omega}(\tilde{\boldsymbol{\omega}}_t)_{(k,h)}}{1+\exp(-\tilde{\omega}_{j,t})}, & \text{if } h < j \quad \wedge \quad j \neq J \\ -\sum_{k=1}^{J-1} \mathcal{J}^{\omega}(\tilde{\boldsymbol{\omega}}_t)_{(k,h)}, & \text{if } j = J, \\ 0, & \text{if } h > j. \end{cases} \tag{13}$$

The score of (1) with respect to the mixture weights parameters follows from Proposition 2.1 and is given by $\nabla^{\omega}(\boldsymbol{\omega}_t | \mathbf{y}_t) = \left(\frac{p_j(\mathbf{y}_t | \boldsymbol{\theta}_{j,t})}{p(\mathbf{y}_t | \boldsymbol{\theta}_t)}, \quad j = 1, \ldots, J \right)'$, i.e. its j–th component is given by the ratio between the mixture distribution and the j–th component distribution. It is interesting to note that, the SDF naturally suggests to update the weights dynamics using the information contained in the density ratio adjusted for the chosen reparametrisation of $\boldsymbol{\omega}_t$, i.e. $\mathcal{J}^{\omega}(\tilde{\boldsymbol{\omega}}_t)' \nabla^{\omega}(\boldsymbol{\omega}_t | \mathbf{y}_t)$. This mechanism implicitly moves the mixture composition to the region of higher probability mass at each

point in time t, and strongly differs from the pure Random Walk assumption implied by Billio et al. (2013).

2.2. Update the mixture components

If we assume no ordering restriction to the mixture distribution, the matrix $\mathcal{J}\left(\tilde{\boldsymbol{\theta}}_t\right)$ turns out to be block diagonal with respect the sub–matrices $\mathcal{J}^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right), \quad j = 1, \ldots, J$, and hence, the dynamic updating equation for $\tilde{\boldsymbol{\theta}}_t$ can be divided into J dynamics for each $\tilde{\boldsymbol{\theta}}_{j,t}$ plus the dynamic for $\tilde{\boldsymbol{\omega}}_t$. This feature of the model also permits to easily parallelise the update of $\tilde{\boldsymbol{\theta}}_t$.[1] As for the update of the mixture composition, the reparametrised mixture component parameters $\tilde{\boldsymbol{\theta}}_{j,t} \quad j = 1, \ldots, J$ are updated using the conditional score of the reparametrised mixture distribution $\tilde{p}\left(\mathbf{y}_t|\tilde{\boldsymbol{\theta}}_t\right)$ with respect to the vector $\tilde{\boldsymbol{\theta}}_{j,t}$, given by $\mathcal{J}^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right)' \nabla^j\left(\boldsymbol{\theta}_{j,t}|\mathbf{y}_t\right)$ where,

$$\nabla^j\left(\boldsymbol{\theta}_{j,t}|\mathbf{y}_t\right) = \omega_{j,t}\frac{p_j\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right)}{p\left(\mathbf{y}_t|\boldsymbol{\theta}_t\right)}\nabla_{p_j}\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right) \quad , j = 1, \ldots, J, \tag{14}$$

follows from Proposition 2.1, and $\nabla_{p_j}\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right) = \frac{\partial \ln p_j\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right)}{\partial \boldsymbol{\theta}_{j,t}}$ is the score of the j–th mixture component. The updating equation for the j–th mixture component parameters can be written as

$$\tilde{\boldsymbol{\theta}}_{j,t+1} = \boldsymbol{\kappa}^j + \xi_{j,t}\mathbf{A}^j\mathcal{J}^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right)' \nabla_{p_j}\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right) + \mathbf{B}^j\tilde{\boldsymbol{\theta}}_{j,t}, \tag{15}$$

where

$$\xi_{j,t} = \omega_{j,t}\frac{p_j\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right)}{p\left(\mathbf{y}_t|\boldsymbol{\theta}_t\right)}, \tag{16}$$

[1]Sometimes, In real applications, it can results to be quite restrictive. Indeed, within the context of (Hidden) Mixture Models, it is often desirable to identify the mixture component parameters using some predefined schemes based on previous knowledge. For example, in the financial markets, during periods of high volatility also the correlations between firms tend to be higher and vice versa. This situation would suggest an ordering of the mixture components variances and correlations. If this is the case, the matrix $\mathcal{J}\left(\tilde{\boldsymbol{\theta}}_t\right)$ is no longer block diagonal and the J dynamics for the mixture components parameters cannot be divided anymore. Clearly, this kind of assumptions can be easily incorporated in DAMM imposing a different parametrisation. For simplicity, in this paper we will ignore this possibility.

is the relative contribution of the j–th component to the mixture density at time t conditionally on past information. It is worth remarking that, equation (15) is very similar to that usually found in Score Driven processes. Indeed, if $J = 1$, we recover the Generalised Autoregressive Score (GAS) and the Dynamic Conditional Score (DCS) models with identity scaling matrix of Creal et al. (2013) and Harvey (2013), respectively. However, in our context, the mixture assumption naturally suggests to scale the score contribution in a way that accounts for the relative importance each mixture component has at time t. Interestingly, a similar result has been found by Bazzi et al. (2014) in their Time Varying Hidden Markov Model. Furthermore, thanks to the use of the conditional score, DAMMs share the same "robustness" properties of classical Score Driven models and can be easily developed starting from that specifications, see e.g. Koopman et al. (2015). It follows that, the DAMM embeds a rational learning mechanism for updating the mixture components parameters. Indeed, we do not need to impose any arbitrary learning mechanism such as the one detailed in Billio et al. (2013), simply because for DAMMs it emerges naturally. This implied learning mechanism ensures that the new information contained in the data is shared across the mixture components proportionally to its relative relevance. Indeed, if we are sure that a new observation has been generated from a particular component of the mixture, DAMM will update the parameters of that particular component leaving the others unchanged simply due to the lack of relevant information. We believe that this point is of primary importance and it is quite often neglected especially for Markow Switching models with dynamic state–dependent densities. For example, in models such as the MS–GARCH of Haas (2006) ,the MS–Copula of Fei et al. (2013), the SGASC of Bernardi and Catania (2015) and the MS–VAR of Krolzig (1997), past observations are treated equally across the latent markovian states. Two notable exceptions are the MS–GARCH specifications of Gray (1996) and Klaassen (2002) where, in order to solve the "path dependence problem", the state dependent GARCH specification is scaled using the predicted and filtered state probabilities, respectively. Even if these ad hoc solutions are developed with a different scope, their consequences are quite similar to those naturally implied by DAMMs.
Concerning the estimation of DAMM, it can be easily performed using the Maximum Likelihood (ML) estimator as detailed in the technical report "Maximum likelihood estimator for generalised

autoregressive score models" by F. Blasques, S.J. Koopman and A. Lucas, available at `http://papers.tinbergen.nl/14029.pdf`, see also Blasques et al. (2015). However, as usually happens when mixtures of distributions are employed, problems of multimodality of the likelihood can arise, hence good starting values and multiple tries are required. We suggest to initialise the algorithm using the estimates delivered by a static Mixture models estimated by the EM algorithm, as suggested by Yu (2012) in a similar context. The supplementary material accompanying this paper reports a simulation study on the finite sample properties of the ML estimator for the class of DAMMs.

3. DAMM specifications

In this Section we report four different DAMM specifications that can be used in empirical works. Specifically, we report two univariate specifications for dynamic mixtures of univariate Gaussian and Student–t distributions as well as two specifications for dynamic mixtures of multivariate Gaussian and multivariate Student–t distributions and their related copulas. It is worth remarking that, the updating of the mixture composition is not strictly related to the choice of the mixture components, so the arguments presented in Subsection 2.1 do not account for the specific model parametric form.

3.1. Univariate DAMM specifications

Since the dynamic updating equation for each reparametrised mixture component vector of parameters $\tilde{\boldsymbol{\theta}}_{j,t}$, only depends on the chosen mapping functions $\Lambda^j(\cdot)$, and the score of the mixture component distribution, we only need to specify the jacobian matrices $\mathcal{J}^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right)$ and the scores vectors $\nabla_{p_j}\left(y_t|\boldsymbol{\theta}_{j,t}\right)$, for $j = 1, \ldots, J$, to fully characterise the updating scheme reported in equation (15). Specifically, if we assume that the univariate observation $y_t \in \Re$ is conditionally distributed according to a mixture of J Gaussian distributions, i.e.

$$y_t \sim \sum_{j=1}^{J} \omega_{j,t} p_j\left(y_t|\boldsymbol{\theta}_{j,t}\right), \tag{17}$$

where $\boldsymbol{\theta}_{j,t} = \left(\mu_{j,t}, \sigma^2_{j,t}\right)'$ contains the mean $\mu_{j,t}$ and variance $\sigma^2_{j,t}$, of the j–th Gaussian density p_j, one can easily reparametrise the vector $\boldsymbol{\theta}_{j,t}$ employing the mapping function

$$\Lambda^j := \begin{cases} \mu_{j,t} & = \mu_{j,t} \\ \sigma^2_{j,t} & = \exp\left(\tilde{\sigma}^2_{j,t}\right) \end{cases} \tag{18}$$

which implies the jacobian matrix

$$\mathcal{J}^j\left(\mu_{j,t}, \tilde{\sigma}^2_{j,t}\right) = \begin{pmatrix} 1 & 0 \\ 0 & \exp\left(\tilde{\sigma}^2_{j,t}\right) \end{pmatrix}. \tag{19}$$

Finally, the score vector for the j–th Gaussian component is given by

$$\nabla_{p_j}\left(y_t|\mu_{j,t}, \sigma^2_{j,t}\right) = \begin{pmatrix} \left(y_t - \mu_{j,t}\right)/\sigma^2_{j,t} \\ \left[\frac{(y_t-\mu_{j,t})^2}{\sigma^2_{j,t}} - 1\right]/2\sigma^2_{j,t} \end{pmatrix}, \tag{20}$$

and the dynamic updating equation for $\tilde{\boldsymbol{\theta}}_{j,t} = \left(\mu_{j,t}, \tilde{\sigma}^2_{j,t}\right)'$, is simply given by

$$\begin{pmatrix} \mu_{j,t+1} \\ \tilde{\sigma}^2_{j,t+1} \end{pmatrix} = \begin{pmatrix} \kappa_{\mu_j} \\ \kappa_{\sigma_j} \end{pmatrix} + \begin{pmatrix} \beta_{\mu_j} & 0 \\ 0 & \beta_{\sigma_j} \end{pmatrix} \begin{pmatrix} \mu_{j,t} \\ \tilde{\sigma}^2_{j,t} \end{pmatrix} \tag{21}$$

$$+ \frac{\omega_{j,t} p_j\left(y_t|\mu_{j,t}, \sigma^2_{j,t}\right)}{\sum_{j=1}^{J} \omega_{j,t} p_j\left(y_t|\boldsymbol{\theta}_{j,t}\right)} \begin{pmatrix} \alpha_{\mu_j} & 0 \\ 0 & \alpha_{\sigma_j} \end{pmatrix} \mathcal{J}^j\left(\mu_{j,t}, \tilde{\sigma}^2_{j,t}\right) \nabla_{p_j}\left(y_t|\mu_{j,t}, \sigma^2_{j,t}\right), \tag{22}$$

where, as stated before, a diagonal structure for the matrices $\mathbf{A}^j$ and $\mathbf{B}^j$ has be imposed.

If we believe that a mixture of Gaussian distributions is not suited for the available time series, we can consider a DAMM specification with Student–t mixture components. In this case, we can assume p_j to be a Student–t distribution with mean $\mu_{j,t}$, scale $\psi_{j,t}$ and shape $\nu_{j,t}$ parameters, and define the vector $\boldsymbol{\theta}_{j,t} = (\mu_{j,t}, \psi_{j,t}, \nu_{j,t})'$, such that $\Lambda^j\left(\tilde{\boldsymbol{\theta}}_{j,t}\right) = \boldsymbol{\theta}_{j,t}$. In this case, if we want that y_t maintains well defined first and second conditional moments, we can define Λ^j equals to

$$\Lambda^j := \begin{cases} \mu_{j,t} & = \mu_{j,t} \\ \psi_{j,t} & = \exp\left(\tilde{\psi}_{j,t}\right) \\ \nu_{j,t} & = \exp\left(\tilde{\nu}_{j,t}\right) + 2 \end{cases}, \tag{23}$$

which implies the jacobian matrix

$$\mathcal{J}^j\left(\mu_{j,t},\tilde{\psi}_{j,t},\tilde{\nu}_{j,t}\right)=\begin{pmatrix}1 & 0 & 0\\ 0 & \exp\left(\tilde{\psi}_{j,t}\right) & 0\\ 0 & 0 & \exp\left(\tilde{\nu}_{j,t}\right)\end{pmatrix}. \tag{24}$$

Finally, the score vector for the j–th Student–t component is given by

$$\nabla_{p_j}\left(y_t|\mu_{j,t},\psi_{j,t},\nu_{j,t}\right)=\begin{pmatrix}\frac{(\nu_{j,t}+1)(y_t-\mu_{j,t})}{2\psi_{j,t}^2+(y_t-\mu_{j,t})^2}\\ -\frac{1}{\psi_{j,t}}+\frac{(\nu_{j,t}+1)(y_t-\mu_{j,t})}{\nu_{j,t}\psi_{j,t}^3+\psi_{j,t}(y_t-\mu_{j,t})}\\ h\left(y_t,\mu_{j,t},\psi_{j,t},\nu_{j,t}\right)\end{pmatrix},$$

where

$$\begin{aligned}h\left(y_t,\mu_{j,t},\psi_{j,t},\nu_{j,t}\right)=&\frac{1}{2}\varpi\left(\frac{\nu_{j,t}+1}{2}\right)-\frac{1}{2}\varpi\left(\frac{\nu_{j,t}}{2}\right)-\frac{\pi}{2\nu_{j,t}}\\ &-\frac{1}{2}\log\left(1+\frac{\left(y_t-\mu_{j,t}\right)^2}{\nu_{j,t}\psi_{j,t}^2}\right)+\frac{\left(\nu_{j,t}+1\right)\left(y_t-\mu_{j,t}\right)^2}{2\nu_{j,t}\left(\psi_{j,t}^2\nu_{j,t}+\left(y_t-\mu_{j,t}\right)^2\right)},\end{aligned}$$

and $\varpi\left(x\right)$ is the digamma function. It is worth noting that, the dynamic features of the model should be tailored to the statistical properties of the considered time series. Indeed, if we do not believe that such a rich parametrisation is required by the data, restrictions on the coefficients that determine the updating of $\boldsymbol{\theta}_{j,t}$ can be imposed. For example, the coefficients α_{ν_j} and β_{ν_j}, for $j=1,\ldots,J$, may be constraint to zero in order to avoid time variation in the mixture components shape parameters.

3.2. Multivariate DAMM specifications

The extension to the multivariate case $\mathbf{y}_t\in\Re^d$ is straightforward. Suppose that $\mathbf{y}_t$ is conditionally distributed according to a mixture of J multivariate Gaussian distributions, with mean vector $\boldsymbol{\mu}_{j,t}$ and covariance matrix $\boldsymbol{\Sigma}_{j,t}$, for all $j=1,\ldots,J$. It is convenient to employ the following decomposition of the covariance matrix

$$\boldsymbol{\Sigma}_{j,t}=\mathbf{D}_{j,t}\mathbf{R}_{j,t}\mathbf{D}_{j,t}, \tag{25}$$

where $\mathbf{D}_{j,t} = \text{diag}(\boldsymbol{\sigma}_{j,t})$ and $\boldsymbol{\sigma}_{j,t} = (\sigma_{j,i,t}, i = 1, \ldots, d)'$, where $\sigma_{j,i,t}$ is the conditional standard deviation of $y_{i,t}$ conditionally on p_j. Let define $\boldsymbol{\rho}_{j,t} = \text{vechd}(\mathbf{R}_{j,t})$ where $\text{vechd}(\mathbf{X})$ is the vech operator without considering the diagonal elements of $\mathbf{X}$. The $(2d + d(d-1)/2)$–valued vector of time varying parameters for the j–th mixture component is $\boldsymbol{\theta}_{j,t} = \left(\boldsymbol{\mu}'_{j,t}, \boldsymbol{\sigma}'_{j,t}, \boldsymbol{\rho}'_{j,t}\right)'$ and its reparametrised version is given by $\tilde{\boldsymbol{\theta}}_{j,t} = \left(\boldsymbol{\mu}'_{j,t}, \tilde{\boldsymbol{\sigma}}'_{j,t}, \tilde{\boldsymbol{\rho}}'_{j,t}\right)'$. Under this parametrisation, the only difficult task is to define a mapping function that maps $\tilde{\boldsymbol{\rho}}_{j,t}$ into $\boldsymbol{\rho}_{j,t}$ such that $\mathbf{R}_{j,t} = \text{vechd}^{-1}\left(\boldsymbol{\rho}_{j,t}\right)$ is a symmetric positive defined correlation matrix, since the exponential mapping function can still be used for the components of $\tilde{\boldsymbol{\sigma}}_{j,t}$. We suggest to use the hyperspherical coordinates reparametrisation for the correlation matrices $\mathbf{R}_{j,t}$, $j = 1, \ldots, J$ as in Creal et al. (2012) and Jaeckel and Rebonato (1999). Specifically, we define $\Lambda^j_\rho : \Re^{d(d-1)/2} \to \Omega^\rho$, such that $\Lambda^j_\rho\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) = \boldsymbol{\rho}_{j,t}$, where $\Omega^\rho = \{(\rho_{j,il,t}, i = 1, \ldots, d, i < l < d) \in [-1,1]^{d(d-1)/2} \mid \Re(\mathbf{x}^* \mathbf{R}_{j,t} \mathbf{x}) > 0, \text{for } \mathbf{R}_{j,t} = \text{vechb}^{-1}\left(\boldsymbol{\rho}_{j,t}\right) \wedge \forall \mathbf{x} \in \mathbb{C}^{d(d-1)/2}\}$[2]. The mapping function $\Lambda^j_\rho(\cdot)$ defines

$$\Lambda^j_\rho\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) = \text{vechd}\left(\mathbf{Z}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)' \mathbf{Z}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)\right), \tag{26}$$

where $\mathbf{Z}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)$ is a $d \times d$ upper—triangular matrix, that is,

$$\mathbf{Z}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) = \begin{pmatrix} 1 & c_{j,12,t} & c_{j,13,t} & \ldots & c_{j,1d,t} \\ 0 & s_{j,12,t} & c_{j,23,t} s_{j,13,t} & \ldots & c_{j,2d,t} s_{j,1k,t} \\ 0 & 0 & 0 & \ldots & c_{j,3d,t} s_{j,2k,t} s_{j,1k,t} \\ 0 & 0 & 0 & \ldots & c_{j,4d,t} s_{j,3k,t} s_{j,2k,t} s_{j,1k,t} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & c_{j,d-1,d,t} \prod_{l=1}^{d-2} s_{j,ld,t} \\ 0 & 0 & 0 & \ldots & \prod_{l=1}^{d-1} s_{j,ld,t} \end{pmatrix}, \tag{27}$$

[2]here $\mathbf{x}^*$ denotes the conjugate transpose of $\mathbf{x}$

where $c_{j,lk,t} = \cos\left(\tilde{\rho}_{j,lk,t}\right)$ and $s_{j,lk,t} = \sin\left(\tilde{\rho}_{j,lk,t}\right)$, where $\tilde{\rho}_{j,lk,t}$ is the (l,k)–th element of $\tilde{\mathbf{R}}_{j,t} = \text{vechd}^{-1}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)$. Consequently, the mapping function for $\tilde{\boldsymbol{\theta}}_{j,t}$ is given by

$$\Lambda^j := \begin{cases} \mu_{j,i,t} & = \mu_{j,i,t}, \quad \forall i = 1,\ldots,d \\ \sigma_{j,i,t} & = \exp\left(\sigma_{j,i,t}\right), \quad \forall i = 1,\ldots,d \\ \boldsymbol{\rho}_{j,t} & = \Lambda^j_{\rho}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right), \end{cases} \tag{28}$$

which implies a jacobian matrix equals to

$$\mathcal{J}^j\left(\boldsymbol{\mu}_{j,t}, \tilde{\boldsymbol{\sigma}}_{j,t}, \tilde{\boldsymbol{\rho}}_{j,t}\right) = \begin{pmatrix} \mathbb{I}_{dd} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathcal{D}\left(\tilde{\boldsymbol{\sigma}}_{j,t}\right) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{J}^j_{\rho}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) \end{pmatrix}, \tag{29}$$

where $\mathcal{D}\left(\tilde{\boldsymbol{\sigma}}_{j,t}\right) = \text{diag}\left(\exp\left(\tilde{\sigma}_{j,i,t}\right), i = 1,\ldots,d\right)$ and $\mathcal{J}^j_{\rho}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)$ is defined in Creal et al. (2012). The score of the j–th component mixture distribution can be partitioned in $\nabla_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \left(\nabla^{\mu}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right)', \nabla^{\sigma}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right)', \nabla^{\rho}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right)'\right)'$, where $\nabla^{\mu}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \left(\nabla^{\mu_i}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right), i = 1,\ldots,d\right)'$, $\nabla^{\sigma}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \left(\nabla^{\sigma_i}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right), i = 1,\ldots,d\right)'$ and $\nabla^{\rho}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \left(\nabla^{\rho_{il}}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right), i = 1,\ldots,d, i < l < d\right)'$ and

$$\nabla^{\mu_i}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \boldsymbol{\iota}_i'\boldsymbol{\Sigma}^{-1}_{j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right) \tag{30}$$

$$\nabla^{\sigma_i}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = -\sigma_{j,i,t} - \frac{1}{2}\left(\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)'\mathbf{K}_{j,i,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)\right) \tag{31}$$

$$\nabla^{\rho_{il}}_{p_j}\left(\boldsymbol{\theta}_{j,t}\right) = \left(\mathbf{v}_{j,t}'\mathbf{R}^{-1}_{j,t}\mathbf{U}_{i,l}\mathbf{R}^{-1}_{j,t}\mathbf{v}_{j,t}\right) - \mathbf{R}^{-1}_{(i,l),j,t}, \tag{32}$$

where $\mathbf{K}_{j,i,t} = -\mathbf{D}^{-1}_{j,t}\boldsymbol{\iota}_i\boldsymbol{\iota}_i'\mathbf{D}^{-1}_{j,t}\mathbf{R}^{-1}_{j,t}\mathbf{D}^{-1}_{j,t} - \mathbf{D}^{-1}_{j,t}\mathbf{R}^{-1}_{j,t}\mathbf{D}^{-1}_{j,t}\boldsymbol{\iota}_i\boldsymbol{\iota}_i'\mathbf{D}^{-1}_{j,t}$, $\mathbf{v}_{j,t} = \mathbf{D}^{-1}_{j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)$, $\boldsymbol{\iota}_i$ is a vector of zeros with 1 at its i–th element and $\mathbf{U}_{i,l}$ is a matrix of zeros except for its (i,l)–th element which is one.

In a multivariate context, the role of the tail dependence implied by the joint distribution is of primary importance, see e.g. McNeil et al. (2015). The multivariate Student–t distribution is usually employed to deal with the tail dependence and the fat tails that usually characterise the data, especially in the financial literature. In out framework, we can assume that $\mathbf{y}_t$ is conditionally distributed according

to a mixture of multivariate Student–t distributions with j–th component given by

$$p_j\left(\mathbf{y}_t|\boldsymbol{\theta}_{j,t}\right) = \frac{\Gamma\left(\frac{\zeta_{j,t}+d}{2}\right)\left[1+\frac{1}{\zeta_{j,t}}\left(\mathbf{y}_t-\boldsymbol{\mu}_{j,t}\right)'\boldsymbol{\Sigma}_{j,t}^{-1}\left(\mathbf{y}_t-\boldsymbol{\mu}_{j,t}\right)\right]^{-\frac{\zeta_{j,t}+d}{2}}}{\Gamma\left(\frac{\zeta_{j,t}}{2}\right)\left(\zeta_{j,t}\pi\right)^{\frac{d}{2}}\left|\boldsymbol{\Sigma}_{j,t}\right|^{1/2}}, \tag{33}$$

where $\boldsymbol{\mu}_{j,t}$ is a vector of location parameters, $\zeta_{j,t}$ is the shape parameter and $\boldsymbol{\Sigma}_{j,t}$ is a proper symmetric positive definite scale matrix. As for the multivariate Gaussian case, it is convenient to decompose $\boldsymbol{\Sigma}_{j,t}$ in

$$\boldsymbol{\Sigma}_{j,t} = \boldsymbol{\Psi}_{j,t}\mathbf{R}_{j,t}\boldsymbol{\Psi}_{j,t}, \tag{34}$$

where $\boldsymbol{\Psi}_{j,t} = \text{diag}\left(\boldsymbol{\psi}'_{j,t}\right)$ is a diagonal matrix containing the individual conditional scale parameters $\boldsymbol{\psi}_{j,t} = (\psi_{j,i,t}, i=1,\ldots,d)'$, and $\mathbf{R}_{j,t}$ is the correlation matrix associated with the j–th mixture component. Consequently, the j–th vector of parameters at time t is given by $\boldsymbol{\theta}_{j,t} = \left(\boldsymbol{\mu}'_{j,t}, \boldsymbol{\psi}'_{j,t}, \boldsymbol{\rho}'_{j,t}, \zeta_{j,t}\right)'$. Following the same arguments of the multivariate Gaussian case, we can define the reparametrised vector of parameters $\tilde{\boldsymbol{\theta}}_{j,t} = \left(\tilde{\boldsymbol{\mu}}'_{j,t}, \tilde{\boldsymbol{\psi}}'_{j,t}, \tilde{\boldsymbol{\rho}}'_{j,t}, \tilde{\zeta}_{j,t}\right)'$, as well as the mapping function

$$\Lambda^j := \begin{cases} \mu_{j,i,t} & = \mu_{j,i,t}, \quad \forall i=1,\ldots,d \\ \psi_{j,i,t} & = \exp\left(\tilde{\psi}_{j,i,t}\right), \quad \forall i=1,\ldots,d \\ \boldsymbol{\rho}_{j,t} & = \Lambda^j_{\rho}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) \\ \zeta_{j,t} & = \exp\left(\tilde{\zeta}_{j,t}\right)+c, \quad c>0 \end{cases}, \tag{35}$$

with associated jacobian matrix given by

$$\mathcal{J}^j\left(\boldsymbol{\mu}_{j,t}, \tilde{\boldsymbol{\psi}}_{j,t}, \tilde{\boldsymbol{\rho}}_{j,t}, \tilde{\zeta}_{j,t}\right) = \begin{pmatrix} \mathbb{I}_{dd} & \mathbf{0} & \mathbf{0} & 0 \\ \mathbf{0} & \mathcal{D}\left(\tilde{\boldsymbol{\theta}}_{j,t}\right) & \mathbf{0} & 0 \\ \mathbf{0} & \mathbf{0} & \mathcal{J}^j_{\rho}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right) & 0 \\ 0 & 0 & 0 & \exp\left(\tilde{\zeta}_{j,t}\right) \end{pmatrix}. \tag{36}$$

Note that, the choice of the scalar c may influences the existence of the moments of the conditional distribution of $\mathbf{y}_t$, indeed, if the existence of the second conditional moment is

required, the condition $c > 2$ should be imposed. The score of the j–th component mixture distribution can be partitioned in $\nabla_{p_j}(\boldsymbol{\theta}_{j,t}) = \left(\nabla^{\mu}_{p_j}(\boldsymbol{\theta}_{j,t})', \nabla^{\sigma}_{p_j}(\boldsymbol{\theta}_{j,t})', \nabla^{\rho}_{p_j}(\boldsymbol{\theta}_{j,t})', \nabla^{\zeta}_{p_j}(\boldsymbol{\theta}_{j,t})\right)'$, where $\nabla^{\mu}_{p_j}(\boldsymbol{\theta}_{j,t}) = \left(\nabla^{\mu_i}_{p_j}(\boldsymbol{\theta}_{j,t}), i = 1,\ldots,d\right)'$, $\nabla^{\sigma}_{p_j}(\boldsymbol{\theta}_{j,t}) = \left(\nabla^{\sigma_i}_{p_j}(\boldsymbol{\theta}_{j,t}), i = 1,\ldots,d\right)'$ and $\nabla^{\rho}_{p_j}(\boldsymbol{\theta}_{j,t}) = \left(\nabla^{\rho_{il}}_{p_j}(\boldsymbol{\theta}_{j,t}), i = 1,\ldots,d, i < l < d\right)'$ and

$$\nabla^{\mu_i}_{p_j}(\boldsymbol{\theta}_{j,t}) = z_{j,t}\boldsymbol{\iota}_i'\boldsymbol{\Sigma}^{-1}_{j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right) \tag{37}$$

$$\nabla^{\sigma_i}_{p_j}(\boldsymbol{\theta}_{j,t}) = -\sigma_{i,j,t} - \frac{z_{j,t}}{2}\left(\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)'\mathbf{K}_{i,j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)\right) \tag{38}$$

$$\nabla^{\rho_{il}}_{p_j}(\boldsymbol{\theta}_{j,t}) = z_{j,t}\left(\mathbf{v}_{j,t}'\mathbf{R}^{-1}_{j,t}\mathbf{U}_{i,l}\mathbf{R}^{-1}_{j,t}\mathbf{v}_{j,t}\right) - \mathbf{R}^{-1}_{(i,l),j,t} \tag{39}$$

$$\nabla^{\zeta}_{p_j}(\boldsymbol{\theta}_{j,t}) = \frac{1}{2}\varpi\left(\frac{\zeta_{j,t}+d}{2}\right) - \frac{1}{2}\varpi\left(\frac{\zeta_{j,t}}{2}\right) - \frac{d}{2\zeta_{j,t}}\left[\log\left(1+\frac{r_{j,t}}{\zeta_{j,t}}\right) - \frac{(\zeta_{j,t}+d)\,r_{j,t}}{(\zeta_{j,t}+r_{j,t})\,\zeta_{j,t}}\right], \tag{40}$$

where $z_{j,t} = \frac{\zeta_{j,t}+d}{(1+r_{j,t})}$, $r_{j,t} = \left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)'\boldsymbol{\Sigma}^{-1}_{j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)$, $\mathbf{v}_{j,t} = \boldsymbol{\Psi}^{-1}_{j,t}\left(\mathbf{y}_t - \boldsymbol{\mu}_{j,t}\right)$, $\mathbf{K}_{i,j,t} = -\boldsymbol{\Psi}^{-1}_{j,t}\boldsymbol{\iota}_i\boldsymbol{\iota}_i'\boldsymbol{\Psi}^{-1}_{j,t}\mathbf{R}^{-1}_{j,t}\boldsymbol{\Psi}^{-1}_{j,t} - \boldsymbol{\Psi}^{-1}_{j,t}\mathbf{R}^{-1}_{j,t}\boldsymbol{\Psi}^{-1}_{j,t}\boldsymbol{\iota}_i\boldsymbol{\iota}_i'\boldsymbol{\Psi}^{-1}_{j,t}$, $\varpi(\cdot)$ is the digamma function and $\boldsymbol{\iota}_i$ and $\mathbf{U}_{i,l}$ are defined under equation (30).

Since the number of parameters for the multivariate DAMM specifications may became very large when d growths, resulting in longer estimation time and higher computational efforts, we can decide to employ a copula specification for the conditional distribution of $\mathbf{y}_t$, and exploit the two step estimation procedure for conditional copulas detailed in Patton (2006). Here, we present the DAMM specification with t–copula mixture components and recover the mixture of Gaussian copulas specification as a special case. Let $\mathbf{u}_t = (u_{i,t}, i = 1,\ldots,d)'$ be the vector collecting the Probability Integral Transformations (PITs) of $y_{i,t}$, $i = 1,\ldots,d$ according to their marginal distributions $f_i\left(y_{i,t}|\boldsymbol{\eta}_{i,t}\right)$, $i = 1,\ldots,d$. In our context, the marginal distributions are freely defined, for example each f_i can be represented by a proper univariate DAMM specification. Exploiting the Sklar (1959)'s Theorem we can define the conditional density distribution of $\mathbf{y}_t$ as the product between the conditional density distribution of the copula $c(\mathbf{u}_t|\boldsymbol{\Delta}_t, \boldsymbol{\zeta})$ and the conditional marginal density distributions. Under the DAMM specification with t–copula mixture components we assume that

$$c(\mathbf{u}_t|\boldsymbol{\Delta}_t, \boldsymbol{\zeta}^{\mathsf{c}}) = \sum_{j=1}^{J}\omega_{j,t}c_T\left(\mathbf{u}_t|\mathbf{R}_{j,t}, \zeta^{\mathsf{c}}_j\right), \tag{41}$$

where $\boldsymbol{\Delta}_t = \left(\text{vechd}\left(\mathbf{R}_{j,t}\right)', j = 1, \ldots, J\right)'$, $\boldsymbol{\zeta}^{\mathsf{c}} = \left(\zeta_j^{\mathsf{c}}, j = 1, \ldots, J\right)$ and $c_T\left(\mathbf{u}_t|\cdot\right)$ is the t–copula density given by

$$c_T\left(\mathbf{u}_t|\mathbf{R}_{j,t}, \zeta_j^{\mathsf{c}}\right) = \frac{\Gamma\left(\frac{\zeta_j^{\mathsf{c}}+d}{2}\right)\Gamma\left(\frac{\zeta_j^{\mathsf{c}}}{2}\right)^{d-1}\left(1+\frac{\mathbf{x}_{j,t}'\mathbf{R}_{j,t}^{-1}\mathbf{x}_{j,t}}{\zeta_{j,t}^{\mathsf{c}}}\right)}{|\mathbf{R}_{j,t}|^{1/2}\Gamma\left(\frac{\zeta^{\mathsf{c}}+1}{2}\right)^{d}\prod_{i=1}^{d}\left(1+\frac{x_{j,i,t}^2}{\zeta_j^{\mathsf{c}}}\right)^{-\frac{\zeta_j^{\mathsf{c}}+d}{2}}}, \tag{42}$$

where $\mathbf{x}_t = (x_{j,i,t}, i = 1, \ldots, d)'$, and $x_{j,i,t} = \mathcal{T}_{\zeta_j^{\mathsf{c}}}^{-1}\left(u_{i,t}\right)$, where $\mathcal{T}_{\zeta_j^{\mathsf{c}}}\left(\cdot\right)$ is the cumulative density function of a standard Student–t distribution with ζ_j^{c} degree of freedom. It is worth remarking that $c\left(\mathbf{u}_t|\cdot\right)$ is a mixture of copulas, and that mixture of copulas are copulas themselves, see e.g. Durante and Sempi (2015). Similarly to the previous multivariate cases, we can define $\boldsymbol{\rho}_{j,t} = \text{vechd}\left(\mathbf{R}_{j,t}\right)$ and, since the only mixture component time–varying vector of parameters is $\boldsymbol{\rho}_{j,t}$, we can set $\boldsymbol{\theta}_{j,t} = \boldsymbol{\rho}_{j,t}$ and consequently employ the mapping function $\Lambda_j^{\rho}\left(\cdot\right)$ defined in equation (26) in order to define $\boldsymbol{\theta}_{j,t} = \Lambda_j^{\rho}\left(\tilde{\boldsymbol{\theta}}_{j,t}\right)$. In this case, the jacobian matrix is given by $\mathcal{J}_{\rho}^{j}\left(\tilde{\boldsymbol{\rho}}_{j,t}\right)$, and the score vector by $\nabla_{p_j}^{\rho}\left(\boldsymbol{\theta}_{j,t}, \zeta_j\right) = \left(\nabla_{p_j}^{\rho_{il}}\left(\boldsymbol{\theta}_{j,t}, \zeta_j\right), i = 1, \ldots, d, i < l < d\right)'$, where

$$\nabla_{p_j}^{\rho_{il}}\left(\boldsymbol{\theta}_{j,t}, \zeta_j\right) = \frac{\zeta_j + d}{\zeta_j + \mathbf{x}_{j,t}'\mathbf{R}_{j,t}^{-1}\mathbf{x}_{j,t}}\left(\mathbf{x}_{j,t}'\mathbf{R}_{j,t}^{-1}\mathbf{U}_{i,l}\mathbf{R}_{j,t}^{-1}\mathbf{x}_{j,t}\right) - \mathbf{R}_{(i,l),j,t}^{-1}, \tag{43}$$

where we note that, for $\zeta_j \to \infty$ we recover the multivariate Gaussian case. It is worth adding that, for this DAMM specification, the j–th mixture components parameter ζ_j^{c} is not allowed to evolve over time since the score of the t–copula density with respect to ζ_j^{c} is not available in closed form and requires a step of numerical integration. This limitation derives from the way the parameter ζ_j^{c} enters into the conditional t–copula density function reported in equation (42).

4. Simulation studies

To demonstrate the flexibility of DAMMs to represent complicated nonlinear dynamics for the means, variances and correlations of random variables, we report four simulation studies. In the first experiment we focus on the ability of DAMMs to approximate the first two conditional moments as well as the conditional distribution of a DMM specification similar to that of Gerlach et al. (2000). In

the second experiment we focus on the correlation dynamic of a bivariate stochastic vector reporting an experiment similar the study conducted by Engle (2002), while in the third experiment we focus on the ability of DAMMs to adapt to changes in the conditional mixture composition. Finally, in the last experiment we focus on the cost, in terms of filtering precision, of misspecification within various DAMM specifications.

4.1. DAMMs as filters for Stochastic DMMs

We aim to approximate the first two conditional moments as well as the conditional distribution generated by the model

$$y_t \sim \omega_{1,t}\mathcal{N}\left(y_t|\mu_{1,t},\sigma_{1,t}^2\right)+\omega_{2,t}\mathcal{N}\left(y_t|\mu_{2,t},\sigma_{2,t}^2\right) \tag{44}$$

where $\omega_{1,t}=\frac{1}{1+\exp(-\tilde{\omega}_t)}$, $\omega_{2,t}=1-\omega_{1,t}$ and $\sigma_{j,t}^2=\exp\left(\tilde{\sigma}_{j,t}^2\right)$, for $j=1,2$, and

$$\begin{aligned}
\tilde{\omega}_{t+1} &= -0.003+0.99\tilde{\omega}_t+\varepsilon_{t+1}^{\omega}, & \varepsilon_{t+1}^{\omega} &\overset{iid}{\sim}\mathcal{N}\left(\varepsilon_{t+1}^{\omega}|0,1.00\right)\\
\mu_{1,t+1} &= 0.09+0.97\mu_{1,t}+\varepsilon_{t+1}^{\mu_1}, & \varepsilon_{t+1}^{\mu_1} &\overset{iid}{\sim}\mathcal{N}\left(\varepsilon_{t+1}^{\mu_1}|0,0.02\right)\\
\tilde{\sigma}_{1,t+1}^2 &= -0.001+0.98\tilde{\sigma}_{1,t}^2+\varepsilon_{t+1}^{\sigma_1}, & \varepsilon_{t+1}^{\sigma_1} &\overset{iid}{\sim}\mathcal{N}\left(\varepsilon_{t+1}^{\sigma_1}|0,0.04\right)\\
\mu_{2,t+1} &= -0.04+0.98\mu_{2,t}+\varepsilon_{t+1}^{\mu_2}, & \varepsilon_{t+1}^{\mu_2} &\overset{iid}{\sim}\mathcal{N}\left(\varepsilon_{t+1}^{\mu_2}|0,0.06\right)\\
\tilde{\sigma}_{2,t+1}^2 &= 0.004+0.99\tilde{\sigma}_{2,t}^2+\varepsilon_{t+1}^{\sigma_2}, & \varepsilon_{t+1}^{\sigma_2} &\overset{iid}{\sim}\mathcal{N}\left(\varepsilon_{t+1}^{\sigma_2}|0,0.08\right),
\end{aligned} \tag{45}$$

where all the innovations are mutually independent for all t. We label this model SDDM. According to (44) we have that the conditional mean $\mathbb{E}\left(y_t|\boldsymbol{\theta}_t\right)$ and the conditional variance $\mathrm{Var}\left(y_t|\boldsymbol{\theta}_t\right)$ of y_t are

$$\begin{aligned}
\mathbb{E}\left(y_t|\boldsymbol{\theta}_t\right) &= \omega_{1,t}\mu_{1,t}+\omega_{2,t}\mu_{2,t}\\
\mathrm{Var}\left(y_t|\boldsymbol{\theta}_t\right) &= \omega_{1,t}\sigma_{1,t}^2+\omega_{2,t}\sigma_{2,t}^2+\omega_{1,t}\mu_{1,t}^2+\omega_{2,t}\mu_{2,t}^2-\left(\omega_{1,t}\mu_{1,t}+\omega_{2,t}\mu_{2,t}\right)^2,
\end{aligned} \tag{46}$$

where $\boldsymbol{\theta}_t=\left(\omega_{j,t},\mu_{j,t},\sigma_{j,t}^2,j=1,2\right)'$. In order to approximate the first two conditional moments of (44) we specify a DAMM specification with $J=2$ Gaussian components. Formally, the approximating model for (44) is

$$y_t \sim \omega_{1,t}\mathcal{N}\left(y_t|\mu_{1,t},\sigma_{1,t}^2\right)+\omega_{2,t}\mathcal{N}\left(y_t|\mu_{2,t},\sigma_{2,t}^2\right), \tag{47}$$

where $\omega_{1,t} = \frac{1}{1+\exp(-\tilde{\omega}_t)}$, $\omega_{2,t} = 1 - \omega_{1,t}$, and $\sigma^2_{j,t} = \exp\left(\tilde{\sigma}^2_{j,t}\right)$, for $j = 1, 2$, and

$$\tilde{\omega}_{t+1} = \kappa_\omega + \alpha_\omega \frac{\exp\left(-\tilde{\omega}_t\right)}{\left(1+\exp\left(-\tilde{\omega}_t\right)\right)^2} \frac{p_1\left(y_t|\mu_{1,t}, \sigma^2_{1,t}\right) - p_2\left(y_t|\mu_{2,t}, \sigma^2_{2,t}\right)}{\omega_{1,t} p_1\left(y_t|\mu_{1,t}, \sigma^2_{1,t}\right) + \omega_{2,t} p_2\left(y_t|\mu_{2,t}, \sigma^2_{2,t}\right)} + \beta_\omega \tilde{\omega}_t \tag{48}$$

where $p_j\left(y_t|\mu_{j,t}, \sigma^2_{j,t}\right)$ is the density of a Gaussian random variable with mean $\mu_{j,t}$ and variance $\sigma^2_{j,t}$, for $j = 1, 2$, evaluated at y_t. The j–th dynamics for the conditional mean and the reparametrised conditional variance are the reported in equation (21), setting $J = 2$.

To perform our experiment we simulate a path of length $T = 10000$ for $\boldsymbol{\theta}_t$ from (45), then, for each $t = 1, \ldots, T$, we simulate from (44) $B = 1000$ pseudo–observations. Finally we estimate on each series of pseudo–observations $\mathbf{y}^{(b)} = \left(y_1^{(b)}, \ldots, y_T^{(b)}\right)$, for $b = 1, \ldots, B$ the two Gaussian components DAMM specification detailed in Subsection 3.1. Then, for each estimated model, we evaluate the implied conditional mean and conditional variance at each point in time $t = 1, \ldots, T$.[3] We benchmark the DAMM with two Gaussian components with the Markov Switching AR(1)–GARCH(1,1) of Haas (2006) (MSARGARCH, henceforth) with two regimes defined by

$$\frac{y_t - \mu_{s,t}}{\sigma_{s,t}} \Big| \left(\mathsf{S}_t = s, \mathcal{F}_{t-1}\right) \sim \mathcal{N}\left(0, 1\right), \tag{49}$$

where

$$\mu_{s,t} = \bar{\mu}_s + \phi_s y_{t-1} \tag{50}$$

$$\sigma^2_{s,t} = \omega_s + \alpha_s z^2_{s,t-1} + \beta_s \sigma^2_{s,t-1}, \tag{51}$$

where $z_{s,t} = y_t - \mu_{s,t}$ and $|\phi_s| < 1, \omega_s > 0, \alpha_s + \beta_s < 1, \alpha_s > 0, \beta_s > 0$ for $s = 1, 2$ are imposed to preserve weak stationarity of the process and to ensure the positiveness of the second conditional moment of y_t. In equation (49) S_t represents an integer valued stochastic variable defined on the discrete space $\{1, 2\}$ that follows a first order Markov chain with transition probability matrix

[3]The first two conditional moments of DAMM have the same formulation of (46). However, contrary to SDMM, in the context of DAMM conditioning on $\boldsymbol{\theta}_t$ is redundant since $\boldsymbol{\theta}_t$ if fully available at time $t-1$ and hence $\boldsymbol{\theta}_t \in \mathcal{F}_{t-1}$.

$\mathbf{Q} = \{\mathsf{q}_{lk}\}$, where $\mathsf{q}_{lk} = \mathbb{P}(\mathsf{S}_t = k|\mathsf{S}_{t-1} = l)$ is the probability that state k is visited at time t given that at time $t-1$ the chain was in state l, and initial probabilities vector $\boldsymbol{\delta} = (\delta_s, s = 1, 2)'$, $\delta_s = \mathbb{P}(S_1 = s)$, i.e., the probability of being in state s at time 1, see e.g. Hamilton (1989) and Dueker (1997). Conditional on $\mathcal{F}_{t-1}$, the distribution of y_t is a mixture of two Gaussian distributions

$$y_t|\mathcal{F}_{t-1} \sim \mathbb{P}(\mathsf{S}_t = 1|\mathcal{F}_{t-1})\,\mathcal{N}\left(y_t|\mu_{1,t}, \sigma^2_{1,t}\right) + \mathbb{P}(\mathsf{S}_t = 2|\mathcal{F}_{t-1})\,\mathcal{N}\left(y_t|\mu_{2,t}, \sigma^2_{2,t}\right), \tag{52}$$

where $\mathbb{P}(\mathsf{S}_t = s|\mathcal{F}_{t-1})$ for $s = 1, 2$, are the predicted probabilities evaluated using the Hamilton filter, see e.g. Frühwirth-Schnatter (2006). Further comparative results are reported in terms of a static mixture model with two Gaussian components estimated on the simulated data using a fixed moving window of $K = 100$ observations, labelled MMR. Specifically, at each $t = K, \ldots, T$ the static mixture model with two Gaussian components is recursively estimated using the observations $\left(y^{(b)}_{t-K+1}, \ldots, y^{(b)}_t\right)$, for $b = 1, \ldots, B$. The implied mean, variance and mixture composition delivered from each static mixture model is stored for each $t = K, \ldots, T$ and $b = 1, \ldots, B$ and compared versus the true quantities generated from the SDMM specification. The first K means, variances and mixture compositions of the MMR specification are not available and are excluded from the comparative analysis. The MMR and the MSARGARCH models are estimated exploiting the Expectation–Maximisation algorithm of Dempster et al. (1977), where for the MSARGARCH model part of the E–step is maximised numerically, see e.g. Zucchini and MacDonald (2009). Comparative results are given in terms of Mean Absolute Error (MAE) and the Mean Square Error (MSE) between the true conditional mean, variance and mixture composition generated from the SDMM specification and those delivered by the DAMM, MSARGARCH and MMR models. To investigate the performance of the DAMM specification to approximate the conditional distribution implied by the SDMM, we also compute the Average Kullback Leibler (AKL) divergence between the true conditional density reported in equation (44) and that estimated by the DAMM, MSARGARCH and MMR models. The AKL divergence for model $m \in \{\text{DAMM, MSARGARCH, MMR}\}$ for the b–th path is given by

$$\text{AKL}^{(b)} = T^{-1} \sum_{t=1}^{T} \int_{\Re} \log \frac{p(y_t|\boldsymbol{\theta}_t)}{p^m\left(y_t|\hat{\boldsymbol{\theta}}^{(b)}_{m,t}\right)} p(y_t|\boldsymbol{\theta}_t)\,\mathrm{d}y_t, \tag{53}$$

where $p\left(y_t|\boldsymbol{\theta}_t\right)$ represents the true density of the SDMM specification reported in equation (44) at time t and $p^m\left(y_t|\hat{\boldsymbol{\theta}}_{m,t}^{(b)}\right)$ represents the density of model m with estimated filtered parameters $\hat{\boldsymbol{\theta}}_{m,t}^{(b)}$ at time t for the b–th simulated series.

Table B.1 reports the median MSE and MAE for $\mathbb{E}\left(y_t|\boldsymbol{\theta}_t\right)$ and $\text{Var}\left(y_t|\boldsymbol{\theta}_t\right)$ as well as the median AKL between the true and estimated conditional distributions for the DAMM, MSARGARCH and MMR models. Comparative results are reported relative to the MSARGARCH. We note that the DAMM with two Gaussian components reports lower MSE and MAE for the conditional mean variance and mixture composition confirming its superior filtering ability versus the MSARGARCH and MMR models. Looking at the last column of Table B.1, we also note that the median AKL divergence between the true and the estimated conditional densities of the DAMM specification is substantially lower than those evaluated according to the MSARGARCH and MMR models. This result suggests that the DAMM specification is more adequate to approximated the conditional densities implied by a SDMM specification and confirms the previous findings of Koopman et al. (2015) for the DAMM specification here considered. To conclude our analysis, in Figure A.1 we report the real as well as the median across the B estimates for the conditional mean, the conditional variance and the mixture weight at each point in time t. We note that, even if the estimated model is highly misspecified, the DAMM is able to accurately represent the dynamic of the first two conditional moments generated from the SDMM reported in the first and second panels of Figure A.1, respectively. Furthermore, also the mixture composition dynamics reported in the third panel of Figure A.1 is accurately approximated by the DAMM specification.

It is worth stressing that, in this MC experiment, DAMM are found to be able to approximate the conditional mean and the conditional variance reported in equation (46), and not the mixture components parameters dynamic reported in equation (45). We like to emphasise that the processes $\{\mu_{j,s}, s > 0\}$ and $\{\tilde{\sigma}_{j,s}, s > 0\}$, $\quad j = 1, 2$ cannot be properly estimated during periods when the mixture composition is highly shifted in favour of a particular component, this because, during those periods, most of the data are generated by the component with higher probability. Consequently, no relevant information is contained in the new data in order to update the parameters of the component

with low weight in the mixture. As stated before, these kind of logical reasonings are quite neglected in the literature and most of the available models do not consider the possibility that, during some periods, they cannot update part of the latent processes. On the contrary, as previously detailed, DAMM naturally incorporates a rationale mechanism that addresses this particular issue.

4.2. Time varying correlations

Our experiment is designed similarly to that of the study conducted by Engle (2002) and Creal et al. (2012). Specifically, we simulate $B = 500$ time series of length $T = 1000$ from a bivariate Gaussian distribution with unit variance and time–varying correlation ρ_t. The dynamics we impose to the correlation parameter are

- Constant: $\rho_t = 0.9$
- Sine: $\rho_t = 0.5 + 0.4 \cos(2\pi t/200)$
- Fast Sine: $\rho_t = 0.5 + 0.4 \cos(2\pi t/20)$
- Step: $\rho_t = 0.9 - 0.5\,(t > 500)$
- Ramp: $\rho_t = \mod(t/200)/200$
- Model1: $\rho_t = \exp(\tilde{\rho}_t) / [1 + \exp(\tilde{\rho}_t)]$ where $\tilde{\rho}_t = -0.4\,(1 - 0.99) + 0.99\tilde{\rho}_{t-1} + 0.14\eta_t^{\rho}, \quad \eta_t^{\rho} \sim \mathcal{N}(0, 1)$
- Model2: $\rho_t = \omega_t \rho_{1,t} + (1 - \omega_t)\,\rho_{2,t}$ where $\rho_{i,t} = \exp(\tilde{\rho}_t) / [1 + \exp(\tilde{\rho}_t)]$, $\tilde{\rho}_{i,t} = \bar{\rho}_i\,(1 - 0.99) + 0.99\tilde{\rho}_{i,t-1} + 0.14\eta_t^{\rho_i}, \quad i = 1, 2,$ and $\bar{\rho}_1 = -0.4,\ \bar{\rho}_2 = 0.4,$ where $\omega_t = [1 + \exp(\tilde{\omega}_t)]^{-1}$, $\tilde{\omega}_t = 0.98\tilde{\omega}_{t-1} + \eta_t^{\omega}$ and $\eta^{\rho_1}, \eta^{\rho_2}, \eta_t^{\omega}$ are iid $\mathcal{N}(0, 1)$.

We estimate by Maximum Likelihood the DAMM with two bivariate Gaussian mixture components to each series of simulated data. Consequently, we evaluate the ability of the model to track the dynamic correlation using the MAE and the MSE between the filtered and the simulated dynamic. We also specify the two constraint DAMM versions DAMM–$\bar{\rho}$ and DAMM–$\bar{\omega}$. The DAMM–$\bar{\rho}$ specification is

defined as the DAMM but with constant state dependent correlations, i.e. $\rho_{j,t} = \rho_j, \quad j = 1, 2$, while the DAMM–$\bar{\omega}$ is again defined as the DAMM but with static mixture composition, i.e. $\omega_t = \omega, \quad \forall t$. For comparative purposes we benchmark the DAMM model with the Dynamic Conditional Correlation (DCC) model of Engle (2002) and to the Exponentially Weighed Moving Average (EWMA) model defined as

$$Q_t = \lambda Q_{t-1} + (1 - \lambda)\, y_{t-1} y_{t-1}' \tag{54}$$

$$\rho_t^{\text{EWMA}} = q_{12,t} / \sqrt{q_{11,t} q_{22,t}}, \tag{55}$$

where $q_{ij,t}$ is the (i, j)–th element of Q_t and λ is set equal to 0.96, Q_1 is fixed as the empirical unconditional correlation of the series. For the DAMM specification we impose the mixture means and variances to be constant and centered at the values $\mu_{j,t} = 0$ and $\sigma_{j,t} = 1$ for $j = 1, 2$ and $t = 1, \ldots, T$. Similar constraints are imposed to the DCC model. Table B.2 reports the median MAE and the median MSE, across the B replications, for the three models and the seven correlation patterns. The results are presented relative to the DCC model. We note that the DAMM specification is preferred versus the considered alternatives in every case. More precisely, the unrestricted DAMM specification is preferred for the Sine, FastSine, Ramp and Model2 cases, the DAMM–$\bar{\rho}$ for the Const and Step, while the DAMM–$\bar{\omega}$ when the Model1 specification is assumed for the evolution of the conditional correlation. Generally, we found the DAMM specifications outperform the DCC and the EWMA models both under the MAE and MSE criterions.

4.3. Time varying mixture composition

Our third simulation experiment focuses on the ability of DAMMs to model the dynamic mixture composition generated according to several patterns. To this end we specify a mixture of two univariate Gaussian distributions with fixed means and variances, i.e. our DGP is of the form

$$y_t \sim \omega_t \mathcal{N}(y_t| - 4, 6) + (1 - \omega_t)\, \mathcal{N}(y_t|1, 3) \tag{56}$$

where ω_t evolves according to one of the following patterns:

- Constant: $\omega_t = 0.9$
- Sine: $\omega_t = \cos(2\pi t/200)$
- Fast Sine: $\omega_t = \cos(2\pi t/20)$
- Step: $\omega_t = 0.9 - 0.5\,(t > 500)$
- Ramp: $\omega_t = \mod(t/100)/100$
- Model1: $\omega_t = [1 + \exp(\tilde{\omega}_t)]^{-1}$ where $\tilde{\omega}_t = -0.015 + 0.98\tilde{\rho}_{t-1} + 0.1\eta_t^{\omega}, \quad \eta_t^{\omega} \sim \mathcal{N}(0,1)$
- Model2: $\omega_t = [1 + \exp(\tilde{\omega}_t)]^{-1}$ where $\tilde{\omega}_t = -0.015 + 0.98\tilde{\rho}_{t-1} + 0.5\eta_t^{\omega}, \quad \eta_t^{\omega} \sim \mathcal{N}(0,1)$.

Model1 and Model2 are nonlinear first order autoregression with different standard deviations assumed for the innovations. According to the selected values of the innovation standard deviations, the ω_t process turns out to evolve more smoothly in the interval $(0, 1)$ for Model1, and to display abrupt changes from 0 to 1 for Model2. We simulate $T = 1000$ observations from (56) assuming that ω_t evolves according to the seven considered patterns. Then we estimate the DAMM with two Gaussian components to the simulated observations and we store the filtered series for ω_t. The procedure is repeated $B = 500$ times. We compare the DAMM filtered dynamics for the mixture composition parameter ω_t with those delivered by a two regime Markov Switching model with time–invariant Gaussian distributions in each regime (MS), and with the MMR model detailed in the first simulation experiment. For all the considered models, the values of the means and the variances of each Gaussian component distribution are fixed to the true values, i.e. for the DAMM specification we impose $\mu_{1,t} = -4$, $\mu_{2,t} = 1$, $\sigma^2_{1,t} = 6$ and $\sigma^2_{2,t} = 3$ for all $t = 1, \ldots, T$. Table B.3 reports the median MAE and the median MSE, across the B replications, for the three models and the seven patterns for ω_t. The results are presented relative to the MS model. Similarly to the previously reported correlation study, we found that the DAMM specification reports superior results for all the patterns assumed for the mixture composition parameter ω_t. The only exception is when ω_t follows the "FastSine" dynamic where the MS model marginally outperfoms the DAMM.

4.4. The costs of a wrong parametrisation for the conditional correlations dynamics

In our last simulation study we investigate the consequences of estimating a misspecified model to the dynamic conditional correlation of a four dimension random variable. Specifically, we specify four Data Generating Processes which only differs for the implied conditional correlation dynamic. The experiment proceeds sampling $B = 500$ series of length T from one of the four considered DGPs, assuming that it is the true one, and then estimating the four DGPs on the simulated data. The experiment is repeated four times assuming at each time a different true DGP between the ones considered. We consider the following general parametric assumption for the conditional distribution of $\mathbf{y}_t \in \Re^4$,

$$\mathbf{y}_t \sim \omega_t \mathcal{N}\left(\mathbf{0}, \mathbf{R}_{1,t}\right) + \left(1 - \omega_t\right) \mathcal{N}\left(\mathbf{0}, \mathbf{R}_{2,t}\right), \tag{57}$$

where $\mathbf{R}_{j,t}, j = 1, 2$ are full correlation matrices. The four DGPs are defined as

- DGP1: ω_t and $\mathbf{R}_{j,t}, \quad \forall i = 1, 2$ are updated using the score of the conditional distribution of $\mathbf{y}_t$ as detailed in Sections 2 and 3.
- DGP2: equal to DGP1 but $\omega_t = \omega$
- DGP3: equal to DGP1 but $\mathbf{R}_{j,t} = \mathbf{R}_j, \quad \forall j = 1, 2$
- DGP4: assumes $\omega_t = \omega$ and $\mathbf{R}_{j,t} = \mathbf{R}_j, \quad \forall j = 1, 2.$

Since the marginal distribution of the j–th component of the random variable $\mathbf{y}_t$ is standard Gaussian, the model implied conditional correlation of $\mathbf{y}_t$ is simply the averages across the correlations of the mixture components with weights ω_t and $(1 - \omega_t)$, and it is indicated with $\mathbf{R}_t^h$, for $h \in \{DGP1, DGP2, DGP3, DGP4\}$. We consider three different sample sizes: small ($T = 500$), medium ($T = 1000$) and moderately large ($T = 2000$). Parameters values are fixed similar to that obtained in the empirical application assuming high persistence for the correlation matrices $\mathbf{R}_{i,t}$ and the mixture composition parameter ω_t. Results are given in terms of the average Frobenius norm of the matrix $\hat{\mathbf{R}}_t^{k|h} - \mathbf{R}_t^h$, where $\hat{\mathbf{R}}_t^{k|h}$ represents the estimated implied correlation matrix at time t assuming the

DGP k, when the true DPG is h, i.e.

$$\bar{f}_{h|k} = T^{-1} \sum_{t=1}^{T} \sqrt{\sum_{l=1}^{4} \sum_{m=1}^{4} \left(\hat{\rho}_{lm,t}^{k|h} - \rho_{lm,t}^{h} \right)^2}, \tag{58}$$

where $\hat{\rho}_{lm,t}^{k|h}$ and $\rho_{lm,t}^{h}$ are the (l,m)–th element of the matrices $\hat{\mathbf{R}}_t^{k|h}$ and $\mathbf{R}_t^h$, respectively. Table B.4 reports the median across the B samples for the average Frobenius norm $\bar{f}_{h|k}$ for all $h, k \in \{DGP1, DGP2, DGP3, DGP4\}$. The rows of Table B.4 indicate the estimated models, while the columns indicate true DGPs from which the series have been generated. In order to facilitate the comparison, all the results are reported relative to the true DGP. We note that, when the true DGP is DGP1, there is no loss in specifying DGP2 when the sample size is small or medium and we have to pay about 5.5% in terms of estimate precision when the sample size is moderately large. The cost became of the order of 13% and 21.4% if we specify DGP3 or DGP4, respectively. Differently, if the true DGP is DGP2 and we estimate DGP1, for moderately large sample sizes, we do not pay the cost of the misspecification in terms of precision estimate of the conditional correlation dynamics. The results changes if we assume either DGP3 or DGP4. In these cases the cost of misspefication increases with the sample size. Specifically, if the true model is DGP3, we pay about 30% if we estimate one of the other considered DGPs. Generally, we found that if the sample size is small, it is better to not assume a highly parametrised model like DGP1, since the uncertainty in model coefficients will influence the resulting filtered correlation dynamics precision. On the contrary, if the sample size is moderately large, and the true DGP is DGP1 or DGP2, the cost of misspecification is low. Finally, estimating the wrong model when the true DGP is DGP3 or DGP4, has a relatively highly impact on the precision the filtered correlation dynamic.

5. Empirical application

In this section we adopt several univariate and multivariate DAMM specifications to estimate and forecast the conditional distribution of a panel of daily financial log returns. We consider four of the most capitalised US firms namely: Apple Inc. (AAPL), Exxon Mobil Corp. (XOM), Wells Fargo

& Co. (WFC) and the General Electric Co. (GE). The data set is obtained from Datastream and consists of 4968 observations spanning from 8th January, 1996 to 22th January, 2016 and including several crisis episodes that affected the US economy like the Savings and Loan (S&L) crisis of the 1990s, the dot–com bubble of 2000–2002 and the Global Finance Crisis (GFC) of 2007-2008. In line with usual financial time series stylised facts, the series display volatility clusters and a heavy tailed unconditional empirical distribution, see e.g. McNeil et al. (2015) for an exhaustive treatment of univariate and multivariate financial returns stylised facts. Furthermore, evidences of time varying second conditional moment as well as time varying correlations are found according to the ARCH LM test and the LMC test (not reported) of Engle (1982) and Tse (2000), respectively.

5.1. In sample analysis

We compare univariate and multivariate DAMM with GARCH/DCC–type of specifications in an in sample analysis. Specifically, we estimate the DAMM with mixtures of J multivariate Gaussian distributions (DAMM(J)–mG), as well as with mixtures of J multivariate Student–t distributions (DAMM(J)–mT). We also include two DAMM copula specifications. The first assumes that the conditional joint distribution of returns at time t is a mixture of J Gaussian copulas (DAMM(J)–copG), while the second assumes a mixture of J Student–t copula (DAMM(J)–copT). The marginal specifications of DAMMs with copula distributions are given by univariate DAMMs with J Student–t distributions (DAMM(J)–uT) components. For all the considered specifications we impose that the mean of each mixture components is constant over time, but changes across the components. The specifications that include univariate or multivariate Student–t distributions (DAMM(J)–mT, DAMM(J)–copT, DAMM(J)–uT) are imposed to have a constant shape parameter that changes across the mixture components. From an unreported analysis we found that, for the considered series, $J = 2$ is preferred for all the DAMM specifications, consequently, for the rest of the analysis we only consider these kind of models avoiding to report $J = 2$ in the model labels. We benchmark the proposed DAMM specifications with the GDCC model of Cappiello et al. (2006) which, similarly to DAMMs, allows for a heterogeneous behaviour of each conditional correlation to past information. As

for the DAMMs we consider multivariate Gaussian (GDCC–mG) and Student–t GDCC (GDCC–mT), as well as GDCC Gaussian copula (GDCC–copG) and Student–t copula (GDCC–copT) specifications (Jondeau and Rockinger, 2006). The marginal specifications of the GDCC–copG and the GDCC–copT specifications are given by univariate GARCH–T models of Bollerslev (1987). All the models are estimated by Maximum Likelihood. The copula models are estimated exploiting the two step Inference Function for Margins (IFM) estimator detailed in Patton (2006). The GDCC–mG model is estimated exploiting the two step estimator detailed in Engle (2002), while the GDCC–mT is estimated maximising the full likelihood. For all the considered models, starting values are chosen by maximising the likelihood of their time–invariant version, such that, for example, for the DAMM(J)–uT we estimate a static mixture of J Student–t distributions. Estimated coefficients are omitted to save space and are reported in the supplementary material accompanying the paper. R and C++ computer codes are available from the author.

5.1.1. Univariate specifications

Here we assess the adequacy of the considered univariate specifications in representing the conditional distribution of the univariate returns series at each point in time. We compare the DAMM–uT and the GARCH–T models that will be employed for the rest of the analysis as the marginal specifications for the copula models, with their analogues Gaussian versions that we label DAMM–uG and GARCH–G, respectively. In order to investigate the dynamic properties of DAMMs, we also include in our univariate analysis constraint versions of the DAMM–uT and DAMM–uG. Specifically, we consider models with time invariant mixture composition ($\bar{\omega}$), models with time varying mixture composition and time invariant mixture components ($\bar{\theta}$), and static mixture models (st). Summarising, our set of univariate models is given by $\mathcal{M}_u$ = {DAMM–uG, DAMM–uG–$\bar{\theta}$, DAMM-uG-$\bar{\omega}$, DAMM-uG–st, DAMM–uT, DAMM–uT–$\bar{\theta}$, DAMM–uT–$\bar{\omega}$, DAMM–uT–st,GARCH–G, GARCH–T}. Comparative results are given in terms of goodness of fit considering the two penalised likelihood criteria AIC and BIC and in terms of adequacy of the models specifications in representing the dynamic features of the series. Table B.5 reports the two information criteria AIC and

BIC and the log likelihood evaluated at its optimum for all the considered univariate specifications. We note that the DAMM–uT specification reports the highest log likelihood and is preferred according to the AIC for AAPL and GE, the GARCH–T model is always preferred according to the BIC except for GE, where the BIC favours the DAMM–uT-$\bar{\omega}$ specification. We also note that, in almost every case, static mixture models (st) as well as constraints models ($\bar{\omega}$, $\bar{\theta}$) are suboptimal compared with their unrestricted counterparts according to both AIC and BIC. Concerning the adequacy of the distributional assumption of each model, we employ the same testing procedure of Diebold et al. (1998), see also Jondeau and Rockinger (2006) and Bernardi and Catania (2015). The procedure tests if the PITs according to the estimated conditional distributions are iid uniformly distributed into the interval (0,1). The "iid" part of the test consists of a Lagrange Multiplier (LM) test on the coefficient of determination of an autoregression of order 20 estimated on the k–th power of the estimated PITs. This test is labelled DGT $-$ AR$^{(k)}$ for $k = 1, 2, 3, 4$, and is distributed according to a $\chi^2(20)$ with a critic value of about 31.4 at the 5% confidence level. The test is is useful to investigate the presence of correlation of the k–th moment of the estimated PITs, see Diebold et al. (1998). Table B.6 reveals the usefulness of accounting for time variation in the mixture components and in the mixture composition. Indeed, we note that the DAMM–uT specification is always able to adequately model the serial dependence that characterises the second and the fourth conditional moments of the considered series. Differently, the GARCH–G and the GARCH–T specifications result to be not able to totally explain the conditional variance of AAPL and WFC. It is interesting to note that, differently from the GARCH cases, where the LM test rejections occur under both the Gaussian and the Student–t assumptions for the conditional distribution of returns, the results of the DAMM–uG and the DAMM–uT are different. Indeed, the LM test for the DAMM–uT specification suggests to accept the null of correct specification of the second and fourth conditional moments for all the considered series. Differently, the LM tests for the DAMM–uG specifications are against the null for the WFC and the GE series. This apparently counterintuitive result is directly linked with the score updating mechanism that moves the conditional variance. Indeed, as detailed in Section 2 and in Creal et al. (2013), the score updating mechanism uses the information coming from the

full conditional distribution, and not only that contained in its expected value as happens in the GARCH case. This means that, if the conditional distributional assumption is appropriate, it is more likely that also the dynamic properties of the series are properly accounted for. To conclude the goodness of fit analysis, the fifth row of table B.6, named DGT − H, reports the test for the uniform assumption of the unconditional distribution of the PITs. We employ again the same testing procedure detailed in Diebold et al. (1998) and Jondeau and Rockinger (2006). The DGT − H test statistic is approximately asymptotically distributed according to a $\chi^2(19)$ distribution with a critical value of about 30.14 at a confidence level of 5%. We note that, the Gaussian distribution performs poorly in describing the conditional distribution of financial returns, while the Student–t is clearly more appropriate. We also found that, mixtures of Gaussian and mixtures of Student–t distributions are generally appropriate for financial returns if the dynamic feature of the considered series are taken into account. Specifically, we found that models with a dynamic mixture composition perform better than models with static mixture composition as it is possible to note comparing, for example, the DAMM–uG and the DAMM–uG–$\bar{\omega}$ specifications.

Figure A.2 reports the estimated conditional standard deviation and the estimated conditional mixture composition implied by the DAMM–uT estimated on each univariate series. The mixture composition is represented by the weights assigned to the second component of the mixture at each point in time t. The conditional standard deviations are those implied by the conditional mixture distribution. We also compare the conditional standard deviations implied by the DAMM–uT with those delivered by the GARCH–T models. For each marginal model, the second mixture component reports a higher coefficient associated to the conditional score that updates the j–th Student–t scale parameter (ψ). Estimated coefficients are not reported to save space and are available in the supplementary material accompanying this paper. These finding implies that the second component scale parameter reacts more to the new information coming from the market. Looking at the time evolution of the mixtures composition we note several interesting results. For example, we found that for all the considered series, during the first part of the sample, the mixture composition is shifted in favour of the second component. The mixtures composition starts to change around the begin of 2004

that coincides with the end of the turbulent period implied by the dot–com bubble of 2000-2002. We note that, for AAPL and XOM, the mixture composition changed sharply, while for WFC and GE the change is more smooth. Concerning the second part of the sample, the mixture composition is more heterogeneous between the considered series. However, for all the considered series, the begin of the GFC implies an abrupt change in the mixture composition in favour of the second component.

5.1.2. Multivariate specifications

Similarly to the univariate analysis we compare multivariate DAMMs with GDCC models. The in sample comparison between multivariate models is reported in terms of goodness of fit considering the AIC and the BIC. The eight specifications we consider are $\mathcal{M}_m$ = {DAMM–mG, DAMM–mT, DAMM–copG, DAMM–copT, GDCC–mG, GDCC–mT, GDCC–copG, GDCC–copT}, where the DAMM–copG and the DAMM–copT share the same univariate specifications which are DAMM–uT, while for the GDCC copula models the univariate specifications are assumed to be GARCH–T. Table B.7 reports the AIC and the BIC for the eight multivariate specifications we consider. Similarly to the univariate analysis, we found that DAMMs report higher log likelihood values, and are always preferred according to the AIC versus their GDCC counterparts. According to the rank induced by the BIC, we found that the DAMM–gCop is preferred versus the GDCC–gCop, and that the DAMM–mG is preferred versus the GDCC–mG. Turning into a global comparison, we found that the DAMM–tCop reports the highest log likelihood and is preferred according to the AIC, while the GDCC-tCop is preferred according to the BIC.

Figure A.3 reports the conditional correlations implied by the DAMM–copT and the GDCC–copT. Conditional correlation are evaluated empirically based on 100'000 simulated draws from the joint conditional distribution implied by the DAMM–tCop and the GDCC–copT specifications as in Chollete et al. (2009). Similarly to the conditional standard deviations, we found that the conditional correlations implied by the two models are quite similar. Notable differences are present in the first part of the sample for the pairs AAPL–WFC, XOM–WFC and AAPL–XOM, and just before the turbulent periods of 2000–2002 and 2007–2008 for the pair XOM–GE. Generally, we found that

correlations between the considered asset have increased during our sample period with a remarkable upward trend concurrently to the crisis periods experienced by the US economy.

5.2. Out of sample analysis

To verify the ability of DAMMs to accurately predict the one step ahead conditional distribution of the considered panel of returns, we report an out of sample analysis based on a scoring rule. To this end, the last $H = 1500$ observations from the 2nd February, 2010 to the end of the sample are regarded as a validating sample and used to compare the considered models. We estimate the univariate and multivariate specifications detailed in the in sample analysis, then we perform a rolling one step ahead forecast of the conditional distribution of returns for the whole validating period. Models parameters are updated each 40 observations (2 months) using a fixed window. Comparative results are given both in terms of univariate and multivariate specifications. Specifically, for each model $m \in (\mathcal{M}_u, \mathcal{M}_m)$, we report the sum of the predictive log scores, i.e.

$$\mathrm{LS}^m = \sum_{s=1}^{H} \log p^m \left(\mathbf{y}_{F+s} | \hat{\boldsymbol{\theta}}^m_{F+s} \right), \tag{59}$$

where F is the length of the in sample period and $p^m \left(\mathbf{y}_{T+s} | \hat{\boldsymbol{\theta}}^m_{T+s} \right)$ is the one step ahead predictive distribution of model m at time $F + s$ conditional on the information at time $F + s - 1$. To further confirm our findings, we also apply the Model Confidence Set (MCS) procedure of Hansen et al. (2011) to the series of negative log scores. [4] Table B.8 reports the log scores for the univariate specifications. We note that for AAPL and GE the model that reports the highest log score is the DAMM–uT, for XOM the best model is DAMM-uG, and for WFC the constraints specification DAMM–uT–$\bar{\omega}$. Concerning the SSM delivered by the MCS procedure, we found that for AAPL only

[4] The MCS procedure is an iterative model selection algorithm that delivers the "Superior Set of Models" (SSM) for which the null hypothesis of Equal Predictive Ability (EPA) is not rejected at a certain confidence level. In our case the "predictive ability" at each point in time $F + s$, for model m, is measured by $\ell^m_{F+s} = \log p^m \left(\mathbf{y}_{F+s} | \hat{\boldsymbol{\theta}}^m_{F+s} \right)$, and the vector of "losses" , used to test the null of EPA at each iteration is given by $\tilde{\boldsymbol{\ell}}^m = (-\ell^m_{F+s}, \quad s = 1, \ldots, H)'$, see Hansen et al. (2011) for further details.

the DAMM-uT specification belongs to the set, for XOM it is given by {DAMM-uG, DAMM-uT, GARCH-T}, for WFC by {DAMM–uT,DAMM–uT-$\bar{\omega}$, GARCH-T} and for GE by {DAMM–uT , DAMM–uT-$\bar{\theta}$, DAMM–uT–$\bar{\omega}$}. Table B.9 reports the LS associated to the multivariate specifications detailed in Subsection 5.1.2. Also in this case we found that DAMM specifications outperform comparative GDCC specifications in terms of predictive accuracy of the one step ahead joint conditional distribution of the considered panel of returns. Indeed, the DAMM–tCop specification reports a log score of about -9222, while the LS associated to the best GDCC specification is only -9289.971. Furthermore, according the MCS test, the SSM for the multivariate specifications is given by {DAMM–gCop, DAMM–Mt, DAMM–tCop} and do not include any GDCC specification. To conclude our analysis we report an out of sample comparison between the DAMM-tCop specification and two of its restricted versions. In this way we want to investigate the benefits, in terms of predictive accuracy of the joint conditional distribution of returns, of having dynamic correlations and dynamic mixture composition. To this end, we introduce the DAMM–tCop–$\bar{\omega}$ and the DAMM–tCop–$\bar{\rho}$ specifications, where the $\bar{\rho}$ and $\bar{\omega}$ labels have the same meaning of Section 4.2. It is worth noting that all the three specifications are able to reproduce time variation in the second conditional moment of $\mathbf{y}_t$, but with different degree of flexibility. Furthermore, since we are modeling the conditional copula distribution of our panel of returns, the mixture weights only influence the dependence structure of the returns series and not the means and the variances. Indeed, the marginal conditional distributions of the series are the same for the three models and do not influence this part of the analysis. The sum of the predictive log scores for the DAMM–tCop–$\bar{\omega}$ and the DAMM–tCop–$\bar{\rho}$ specifications are -9281.876 and -9243.246, respectively. Recalling that the sum of the predictive log scores for the DAMM–tCop specification is -9222.006, we found an improvement of about 60 and 21 points in terms of log scores for the DAMM–tCop versus the DAMM–tCop–$\bar{\omega}$ and the DAMM–tCop–$\bar{\rho}$ specifications. To further investigate the differences between the considered models, in Figure A.4 we report the one step ahead predicted correlations implied by the three DAMM copula specifications. We note that the correlations predicted by the DAMM–tCop–$\bar{\omega}$ *(red lines)* are substantially different from those predicted by the DAMM–tCop *(black lines)* and DAMM–tCop–$\bar{\rho}$ *(blue lines)* specifications.

Interestingly, we found that the conditional correlations for the pair XOM–GE are very similar between those implied by the DAMM–tCop and the DAMM–tCop–$\bar{\rho}$ specifications. This is also true for the pair AAPL–XOM during the subperiod 2012/01–20013/12 and for the WFC–GE pair during the last part of the sample. The correlations for the pair XOM–WFC predicted by the DAMM–tCop–$\bar{\rho}$ specification are similar of that implied by the DAMM–tCop but display much lower variability. In all other cases the two conditional correlations are substantially different between the DAMM–tCop and the DAMM–tCop–$\bar{\rho}$ specifications. These results suggest that, the benefits of having time–variation in the mixture composition as well as in the mixture component correlation matrices, are higher in a fully multivariate framework, where more heterogeneity affects for the dependence structure of returns. Indeed, in our analysis with a panel of four returns, we found that for the majority of the cases, allowing for both time–varying mixture composition and mixture component correlation matrices, results in different predictions in terms of implied conditional correlation of the returns. Probably, in a bivariate setting, the benefits of having time–varying mixture composition and time–varying mixture component correlation matrices would be much much lower in terms of flexibility of the implied conditional correlation dynamics. However, the prediction of other characteristics of the conditional distribution, as for example the tail dependence of the data during periods of market turmoil, can still benefit from a richer model parametrisation.

6. Conclusion

In this paper we introduced a new class of Dynamic Mixture Models, named Dynamic Adaptive Mixture Models (DAMMs). DAMMs are flexible tools for mixtures of distributions that dynamically update their composition as well as their components. The updating mechanism is based on the score of the conditional mixture distribution exploiting the recent advances for Score Driven models of Harvey (2013) and Creal et al. (2013). DAMMs are first reparametrised in terms of auxiliary dynamic parameters using an adequate mapping function, then the processes for the auxiliary parameters are updated using the score of the reparametrised conditional mixture at each point in time. Finally,

the parameters of interest are found by mapping the auxiliary parameters into the proper parameter space. DAMMs belong to the class of observation driven models (Cox et al., 1981) with the usual consequence of having the likelihood available in closed form. Hence, DAMMs can be easily estimated by Maximum Likelihood, see e.g. Blasques et al. (2015). A Monte Carlo experiment revels that, the new proposed specification is able to adequately estimate the first two conditional moments as well as the dynamic mixture composition generated from highly nonlinear Stochastic Dynamic Mixture Models as well as several processes assumed for the conditional correlation and the mixture composition of conditional random variables. In the last part of the paper, we report an empirical application in financial econometrics where several univariate and multivariate DAMM specifications are estimated using a panel of financial returns. We found that DAMMs are usually preferred versus GARCH/DCC alternatives both in terms of in sample goodness of fit, and out of sample forecast ability of the marginal and joint conditional distribution. Indeed, we also found that, contrary to the competing models, the DAMM specifications always belong to the Superior Set of Models delivered by the Model Confidence Set of Hansen et al. (2011) under negative Log–Score losses. Furthermore, this holds for very high confidence levels. To conclude, we found DAMM very flexible and easy to handle, we believe they could be successfully employed in other relevant scientific applications such as graphical engineering, biology and spatio–temporal econometrics, where more complicated alternatives are usually employed.

Acknowledgments

I would like to express my sincere thanks to Prof. Andrew Harvey, Prof. Tommaso Proietti and Prof. Siem Jan Koopman, for their constructive comments that greatly contributed to improve the final version of the paper.

Appendix A. Figures

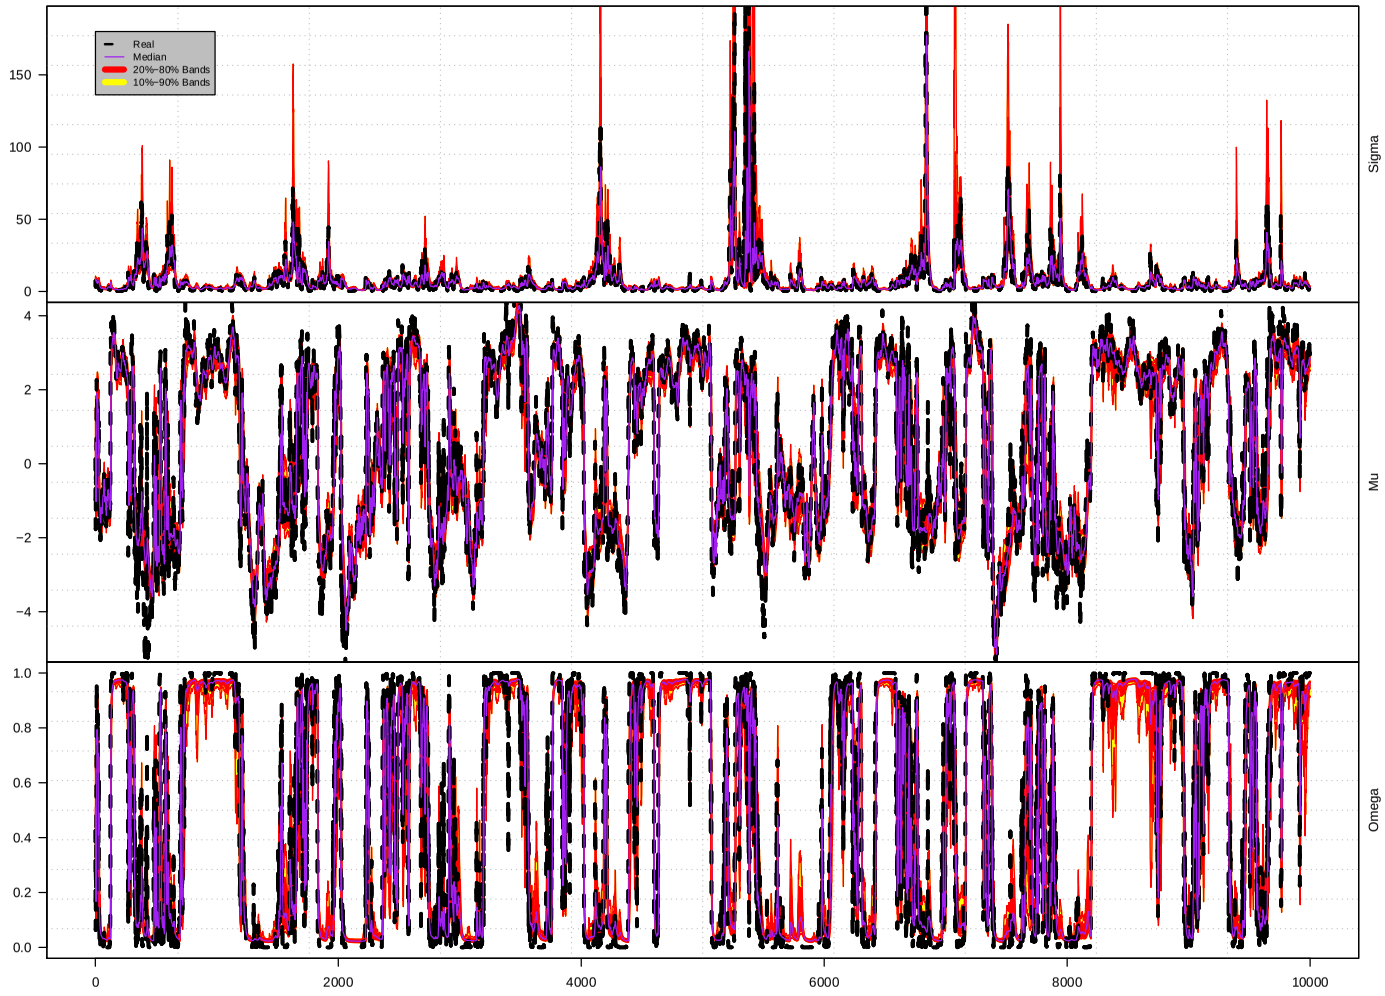

Figure A.1: SDMM approximation using DAMM. Black dotted lines represent the paths for the true conditional standard deviation ("Sigma"), mean ("Mu") and mixture composition ("Omega"). Purple lines are the medians across the 1000 estimates delivered by the DAMM using data simulated from (44) accordingly to the previously simulated paths. Red and yellow bands are 20%–80% and 10%–90% quantiles evaluated at each point in time t using the 1000 estimates, respectively.

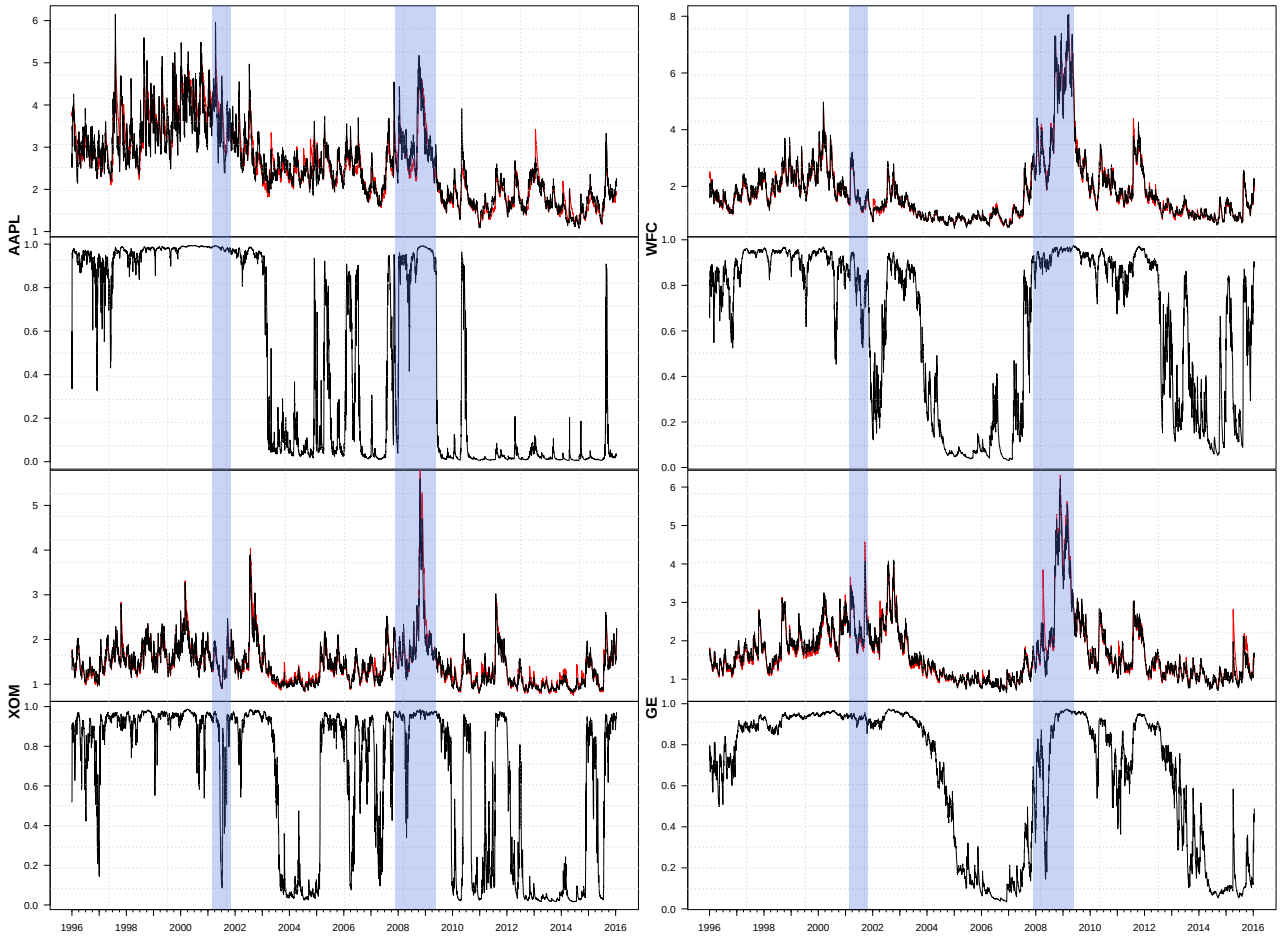

Figure A.2: Conditional standard deviations and mixture composition for the DAMM–uT specification (*black lines*) and conditional standard deviations delivered by the GARCH–T models (*red lines*) for the considered panel of returns. The mixture composition is represented by the weight assigned to the second component of the mixture at each point in time t and it is reported in the bottom figure relative to each return series. The conditional standard deviations are those implied from the conditional mixture distribution and are reported in the top figures relative to each return series. Blue vertical bands represent periods of recessions according to the Recession Indicators Series available from the Federal Reserve Bank of St. Louis

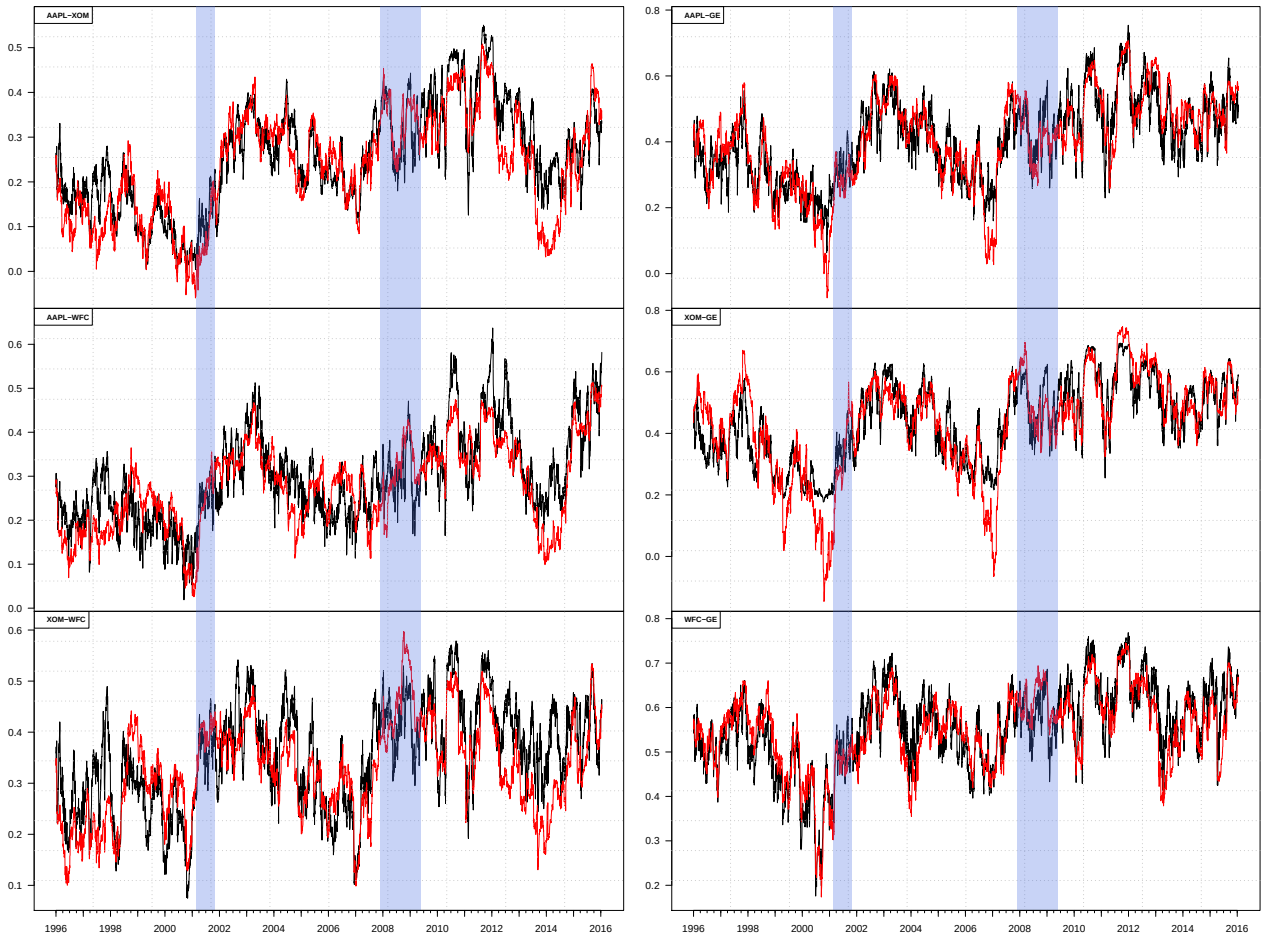

Figure A.3: Conditional linear correlation implied by the DAMM–tCop specification (*black lines*) and GDCC–tCop model (*red lines*) for each pair of returns. For both models the correlations are evaluated empirically using 100'000 samples from the conditional joint distribution at each point in time t. Blue vertical bands represent periods of recessions according to the Recession Indicators Series available from the Federal Reserve Bank of St. Louis

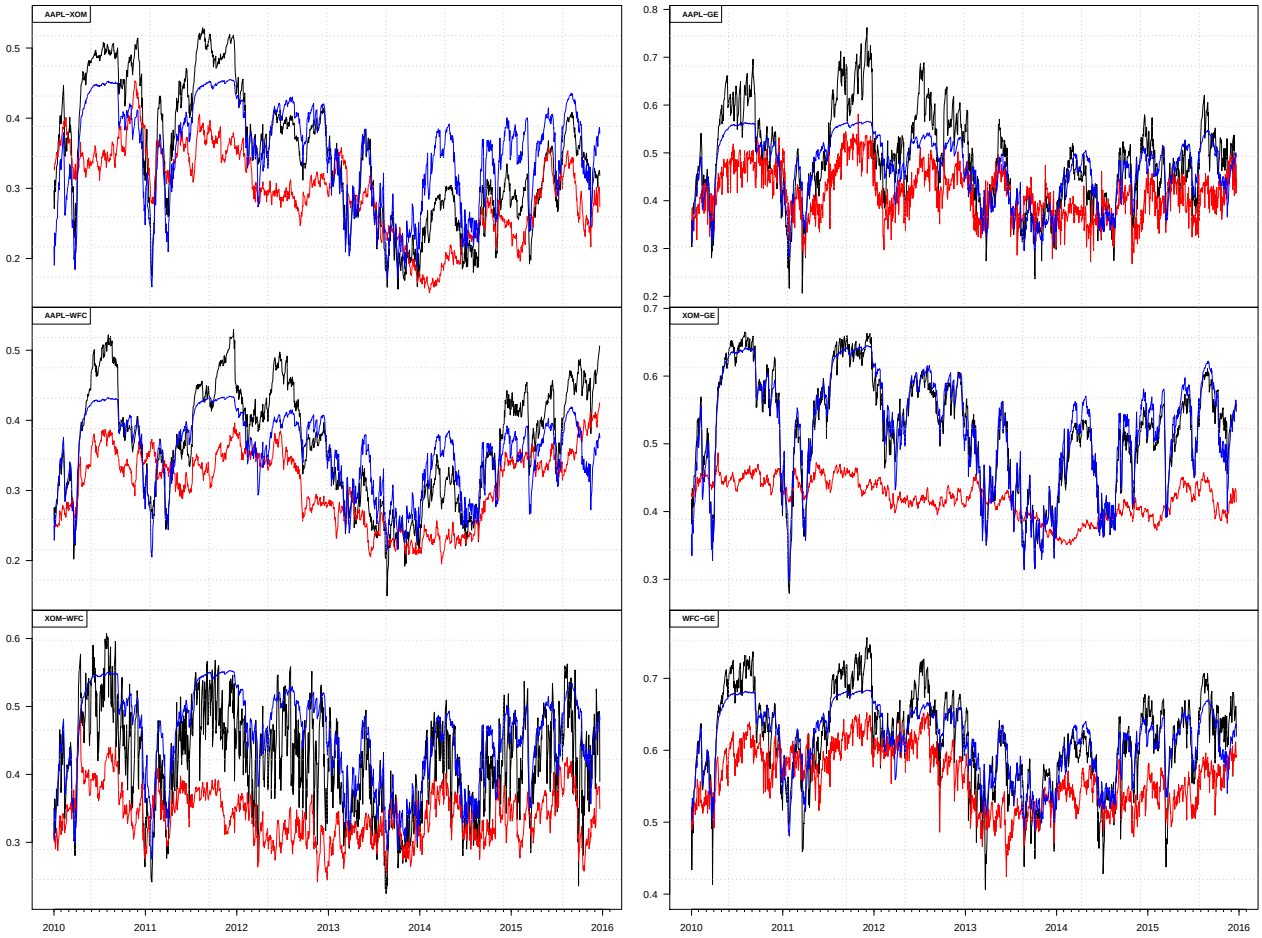

Figure A.4: Rolling one step ahead prediction of the conditional linear correlation implied by the DAMM–tCop (*black lines*) the DAMM–tCop–$\bar{\omega}$ (*red lines*) and DAMM–tCop–$\bar{\rho}$ (*blue lines*) specifications. The correlations are evaluated empirically using 100'000 samples from the predictive conditional joint distribution at each point in time during the validating period

Appendix B. Tables

Spec	MAE			MSE			AKL
	ω_t	$\mathbb{E}(y_t\|\boldsymbol{\theta}_t)$	$\text{Var}(y_t\|\boldsymbol{\theta}_t)$	ω_t	$\mathbb{E}(y_t\|\boldsymbol{\theta}_t)$	$\text{Var}(y_t\|\boldsymbol{\theta}_t)$	
DAMM	**0.196**	**0.587**	**0.737**	**0.038**	**0.345**	**0.543**	**0.637**
MSARGARCH	1.000	1.000	1.000	1.000	1.000	1.000	1.000
MMR	1.277	1.374	1.55	1.631	1.887	2.401	1.412

Table B.1: MAE and MSE between the conditional mean, variance and mixture composition generated by the SDMM specification and estimated from the DAMM, MSARGARCH and MMR models. The last column reports the Average Kullback Leibler (AKL) divergence between the conditional distribution of the SDMM specification and the DAMM, MSARGARCH and MMR models. The reported values are the median across the B replications. Results are presented relative to the MSARGARCH model.

Spec	Const	Sine	FastSine	Step	Ramp	Model1	Model2
MAE							
DAMM	0.9278	**0.9068**	**0.9418**	0.3481	**0.8606**	0.9712	**0.9214**
DAMM–$\bar{\rho}$	**0.8005**	0.9734	0.9428	**0.2934**	1.0859	1.0210	0.9650
DAMM–$\bar{\omega}$	0.8077	0.9884	1.0187	0.7920	1.0184	**0.9532**	0.9495
DCC	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
EWMA	3.9346	1.2847	1.1871	0.9959	1.1100	1.0633	1.0528
MSE							
DAMM	1.1599	**0.8353**	**0.9249**	0.2681	**0.7543**	0.9448	**0.8789**
DAMM–$\bar{\rho}$	**0.8223**	0.9159	0.9224	**0.2293**	1.0699	1.0090	0.9385
DAMM–$\bar{\omega}$	0.7075	0.9241	1.0211	0.6887	0.9884	**0.9106**	0.8951
DCC	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
EWMA	15.3404	1.4543	1.4022	1.0287	1.2867	1.1811	1.1754

Table B.2: MAE and MSE for the estimated conditional correlation dynamic for the DAMM, DCC and EWMA models for seven different patterns. The reported values are the median across the B replications. The results are presented relative to the DCC model.

Spec	Const	Sine	FastSine	Step	Ramp	Model1	Model2
MAE							
DAMM	**0.9951**	**0.84**	1.0005	**0.2979**	**0.7094**	**0.8832**	**0.6451**
MS	1.0000	1.0000	**1.0000**	1.0000	1.0000	1.0000	1.0000
MMR	2.7178	1.4333	1.317	0.289	1.246	0.9727	1.0816
MSE							
DAMM	**0.5232**	**0.7966**	1.0037	**0.1566**	**0.6376**	**0.7966**	**0.5091**
MS	1.0000	1.0000	**1.0000**	1.0000	1.0000	1.0000	1.0000
MMR	4.0257	1.7442	1.5615	0.212	1.4418	0.982	1.2349

Table B.3: MAE and MSE for the estimated ω_t dynamic for the DAMM, MS and MMR models for seven different patterns. The reported values are the median across the B replications. The results are presented relative to the MS model.

	DGP1			DGP2		
T	500	1000	2000	500	1000	2000
DGP1	1.000	1.000	1.000	1.155	1.133	0.965
DGP2	0.916	0.988	1.055	1.000	1.000	1.000
DGP3	1.017	1.085	1.130	1.191	1.292	1.161
DGP4	1.043	1.152	1.214	1.176	1.373	1.290
	DGP3			DGP4		
T	500	1000	2000	500	1000	2000
DGP1	1.248	1.234	1.292	1.481	1.502	1.486
DGP2	1.116	1.143	1.308	1.322	1.352	1.339
DGP3	1.000	1.000	1.000	1.102	1.127	1.125
DGP4	0.884	0.993	1.309	1.000	1.000	1.000

Table B.4: Medians across the B samples of the average Frobenius norm reported in Equation (58) for different sample sizes T. The rows indicates estimated models while the columns indicate true DGPs. All the results are relative to the true DGP.

	DAMM–uG				DAMM–uT				GARCH–G	GARCH–T
	–	$\bar{\theta}$	$\bar{\omega}$	st	–	$\bar{\theta}$	$\bar{\omega}$	st		
					AAPL					
AIC	22862.76	22854.76	23519.24	23512.53	**22712.42**	22841.8	22730.53	23487.96	23042.81	22751.54
BIC	22934.38	22900.34	23577.84	23545.08	22797.06	22900.4	22802.15	23533.54	23068.85	**22784.09**
LLK	-11420.38	-11420.38	-11750.62	-11751.26	**−11343.21**	-11411.9	-11354.27	-11736.98	-11517.4	-11370.77
Np	11	7	9	5	13	9	11	7	4	5
					XOM					
AIC	17202.44	17421.14	17219.17	17836.84	**17182.6**	17350.77	17188.73	17797.33	17299.23	17185.27
BIC	17274.06	17466.72	17277.77	17869.4	17267.24	17409.37	17260.35	17842.9	17325.27	**17217.82**
LLK	-8590.22	-8703.57	-8600.58	-8913.42	**−8578.3**	-8666.39	-8583.37	-8891.66	-8645.62	-8587.63
Np	11	7	9	5	13	9	11	7	4	5
					WFC					
AIC	18446.19	18955.97	20188.08	20011.81	18398.75	18979.92	18393.5	19942.13	18534.78	**18391.85**
BIC	18517.81	19001.55	20246.68	20044.36	18496.41	19038.51	18465.12	19987.71	18560.82	**18424.41**
LLK	-9212.09	-9470.99	-10085.04	-10000.9	**−9184.37**	-9480.96	-9185.75	-9964.06	-9263.39	-9190.93
Np	11	7	9	5	15	9	11	7	4	5
					GE					
AIC	18305.78	18535.34	19442.95	19434.97	18200.29	18445.04	**18198.28**	19243.02	18515.85	18246.59
BIC	18377.39	18580.92	19501.55	19467.53	18284.93	18503.64	**18269.9**	19288.59	18541.9	18279.15
LLK	-9141.89	-9260.67	-9712.48	-9712.49	**−9087.14**	-9213.52	-9088.14	-9614.51	-9253.93	-9118.3
Np	11	7	9	5	13	9	11	7	4	5

Table B.5: AIC, BIC, Log Likelihood (LLK) and number of parameters (Np) for the DAMM-uG, DAMM-uT and related constrained versions and GARCH-G and GARCH-T specifications for each returns series.

	DAMM–uG				DAMM–uT				GARCH–G	GARCH–T
	–	$\bar{\theta}$	$\bar{\omega}$	st	–	$\bar{\theta}$	$\bar{\omega}$	st		
					AAPL					
DGT – AR^1	23.92	23.93	27.31	27.32	24.28	24.25	23.64	27.2	24.32	23.87
DGT – AR^2	70.4^a	70.45^a	678.45^a	678.77^a	18.96	80.44^a	21.01	675.7^a	43.25^a	39.01^a
DGT – AR^3	27.88	27.91	40.32^a	40.37^a	23.58	29.15^c	23.32	39.91^a	24.15	24.5
DGT – AR^4	79.27^a	79.42^a	677.19^a	677.21^a	22.93	94.52^a	27.73	676.15^a	46.91^a	42.97^a
DGT-H	36.25^a	35.94^b	29.63^c	29.08^c	27.56^c	28.22^c	36.77^a	24.04	108.95^a	31.77^b
					XOM					
DGT – AR^1	42.56^a	44.03^a	41.33^a	55.41^a	42.16^a	45.23^a	43.19^a	55.21^a	41.41^a	41.92^a
DGT – AR^2	14.11	50.59^a	16.03	459.24^a	14.41	59.14^a	12.63	457.98^a	14.29	14.18
DGT – AR^3	29.84^c	36.18^b	29.45^c	59.47^a	29.31^c	37.13^b	29.27^c	59.55^a	28.79^c	30.2^c
DGT – AR^4	18.77	62.54^a	20.52	511.55^a	18.46	85.15^a	16.54	515.26^a	19.95	19.05
DGT-H	23.82	19.79	28.67^c	23.2	17.19	19.62	18.01	10.35	52.53^a	16.68
					WFC					
DGT – AR^1	47.92^a	51.44^a	51.88^a	51.61^a	48.86^a	50.53^a	48.73^a	51.24^a	49.38^a	49.35^a
DGT – AR^2	38.1^a	200.77^a	1234.35^a	1181.86^a	20.61	214.99^a	20.04	1187.31^a	35.82^b	29.8^c
DGT – AR^3	37.2^b	41.22^a	59^a	53.6^a	38.58^a	41.45^a	38.4^a	53.6^a	35.64^b	37.19^b
DGT – AR^4	33.3^b	236.24^a	1358.73^a	1301.36^a	18.46	252.61^a	17.34	1311.01^a	30.57^c	24.54
DGT-H	26.75	35.52^b	85.92^a	38.3^a	19.35	69.89^a	20.78	12.98	64.49^a	18.57
					GE					
DGT – AR^1	23.92	22.17	24.21	24.21	24.08	22.12	23.05	24.22	22.27	22.76
DGT – AR^2	20.61	109.93^a	853.52^a	853.52^a	14.16	105.25^a	15.89	829.29^a	25.73	27.29
DGT – AR^3	18.46	17.82	24.64	24.63	19.34	18.09	19.03	24.51	16.74	17.05
DGT – AR^4	27.24	147.07^a	959.28^a	959.38^a	19.51	154.86^a	22.69	926.86^a	28.06	28.61^c
DGT-H	22.46	26.25	64.17^a	64.04^a	18.25	22.95	11.77	14.75	64.23^a	11.62

Table B.6: In sample Goodness–of–Fit test of Diebold et al. (1998). The apexes "a", "b" and "c", denote the rejection of the null hypothesis of not significance of the corresponding parameter, at different confidence levels 1%, 5% and 10%. See also Jondeau and Rockinger (2006).

Specification	AIC	BIC	LLK	Np
DAMM-tCop	**71954.613**	72560.115	**−35884.307**	93
DAMM-gCop	71970.805	72563.285	-35894.403	91
DAMM-Mt	72087.972	72563.259	-35970.986	73
DAMM-Mg	72427.49	72889.754	-36142.745	71
GDCC-tCop	72069.748	**72297.625**	-35999.874	35
GDCC-gCop	72256.073	72477.439	-36094.037	34
GDCC-Mt	72112.263	72314.097	-36025.131	31
GDCC-Mg	73310.633	73505.956	-36625.317	30

Table B.7: AIC, BIC, Log Likelihood (LLK) and number of parameters (Np) for the multivariate DAMM and the GDCC specifications.

Specification	AAPL	XOM	WFC	GE
DAMM-uG	-2866.035	**−2275.761***	-2618.492	-2509.159
DAMM-uG-$\bar{\theta}$	-2949.23	-2322.903	-2643.335	-2511.027
DAMM-uG-$\bar{\omega}$	-3052.224	-2364.859	-2689.309	-2577.902
DAMM-uG-st	-3061.479	-2419.254	-2778.797	-2618.739
DAMM-uT	**−2812.251***	-2276.699*	-2605.368*	**−2490.477***
DAMM-uT-$\bar{\theta}$	-2862.802	-2303.781	-2654.574	-2496.078*
DAMM-uT-$\bar{\omega}$	-2833.06	-2281.131	**−2603.959***	-2492.675*
DAMM-uT-st	-3052.207	-2411.813	-2767.891	-2605.363
GARCH-G	-2899.621	-2307.692	-2635.302	-2567.939
GARCH-T	-2821.322	-2275.877*	-2606.006*	-2500.099

Table B.8: Out of sample log score for all the considered univariate specifications. For each asset, the model that reports the highest log score is indicated in bold. Asterisks represent models that belong to the Superior Set of Models delivered by the MCS procedure with a probability higher then 95%.

	DAMM-Mg	DAMM-gCop	DAMM-Mt	DAMM-tCop	GDCC-gCop	GDCC-Mg	GDCC-Mt	GDCC-tCop
LS	-9321.838	-9228.399*	-9224.843*	**−9222.006***	-9289.971	-9512.197	-9346.559	-9350.887

Table B.9: Out of sample log score for all the considered multivariate specifications. The model that reports the highest log score is indicated in bold. Asterisks represent models that belong to the Superior Set of Models delivered by the MCS procedure with a probability higher then 95%

References

Ardia, D. (2008). *Financial risk management with Bayesian estimation of GARCH models*. Springer.

Bazzi, M., Blasques, F., Koopman, S. J., and Lucas, A. (2014). Time varying transition probabilities for markov regime switching models.

Bernardi, M. and Catania, L. (2015). Switching–GAS Copula Models for Systemic Risk Assessment. *1504.03733*.

Billio, M., Casarin, R., Ravazzolo, F., and van Dijk, H. K. (2012). Combination schemes for turning point predictions. *The Quarterly Review of Economics and Finance*, 52(4):402–412.

Billio, M., Casarin, R., Ravazzolo, F., and Van Dijk, H. K. (2013). Time-varying combinations of predictive densities using nonlinear filtering. *Journal of Econometrics*, 177(2):213–232.

Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer.

Blasques, F., Koopman, S. J., Łasak, K., and Lucas, A. (2015). In–sample confidence bounds and out–of–sample forecast bands for time–varying parameters in observation driven models. *Tinbergen Institue Discussion Paper*.

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *The review of economics and statistics*, pages 542–547.

Cappé, O., Moulines, E., and Rydén, T. (2005). *Inference in hidden Markov models*, volume 6. Springer.

Cappiello, L., Engle, R. F., and Sheppard, K. (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial econometrics*, 4(4):537–572.

Casarin, R., Grassi, S., Ravazzolo, F., and van Dijk, H. K. (2015). Dynamic predictive density combinations for large data sets in economics and finance.

Chollete, L., Heinen, A., and Valdesogo, A. (2009). Modeling international financial returns with a multivariate regime-switching copula. *Journal of Financial Econometrics*, 7(4):437–480.

Cox, D. R., Gudmundsson, G., Lindgren, G., Bondesson, L., Harsaae, E., Laake, P., Juselius, K., and Lauritzen, S. L. (1981). Statistical analysis of time series: Some recent developments [with discussion and reply]. *Scandinavian Journal of Statistics*, pages 93–115.

Creal, D., Koopman, S. J., and Lucas, A. (2012). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business & Economic Statistics*.

Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795.

Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, 39(1):1 – 38.

Diebold, F. X., Gunther, T. A., and Tay, A. a. (1998). Evaluating density forecasts: With applications to financial risk management. *International Economic Review*, 39(4):863–883.

Dueker, M. J. (1997). Markov switching in garch processes and mean-reverting stock-market volatility. *Journal of Business & Economic Statistics*, 15(1):26–34.

Durante, F. and Sempi, C. (2015). *Principles of copula theory*. CRC Press.

Engle, R. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007.

Fei, F., Fuertes, A.-M., and Kalotychou, E. (2013). Modeling dependence in cds and equity markets: Dynamic copula with markov-switching. *Working Paper*.

Frigessi, A., Haug, O., and Rue, H. (2002). A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes*, 5(3):219–235.

Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models: Modeling and Applications to Random Processes*. Springer.

Gerlach, R., Carter, C., and Kohn, R. (2000). Efficient bayesian inference for dynamic mixture models. *Journal of the American Statistical Association*, 95(451):819–828.

Gray, S. F. (1996). Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27 – 62.

Haas, M. M. (2006). S. and paolella, m. *Modelling and predicting market risk with laplace–Gaussian mixture distributions, Applied Financial Economics*, 16:1145–1162.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 77:357–384.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.

Harrison, J. and West, M. (1999). *Bayesian Forecasting & Dynamic Models*. Springer.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge University Press.

Jaeckel, P. and Rebonato, R. (1999). The most general methodology for creating a valid correlation matrix for risk management and option pricing purposes. *Journal of risk*, 2(2):17–28.

Jondeau, E. and Rockinger, M. (2006). The copula-garch model of conditional dependencies: An international stock market application. *Journal of international money and finance*, 25(5):827–853.

KaewTraKulPong, P. and Bowden, R. (2002). An improved adaptive background mixture model for real-time tracking with shadow detection. In *Video-based surveillance systems*, pages 135–144. Springer.

Kim, C.-J. (1994). Dynamic linear models with markov-switching. *Journal of Econometrics*, 60(1):1 – 22.

Kim, C.-J. and Nelson, C. R. (1999). *State-space models with regime switching: classical and Gibbs-sampling approaches with applications*, volume 2. MIT press Cambridge.

Klaassen, F. (2002). Improving garch volatility forecasts with regime-switching garch. In *Advances in Markov-Switching Models*, pages 223–254. Springer.

Koopman, S. J., Lucas, A., and Scharth, M. (2015). Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics, forthcoming*.

Krolzig, H.-M. (1997). Markov-switching vector autoregressions–modelling. *Statistical Inference, and Application to Business Cycle Analysis*.

McLachlan, G. and Peel, D. (2000). *Finite Mixture Models*. Wiley Series in Probability and Statistics, John Wiley & Sons, New York.

McNeil, A. J., Frey, R., and Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools: Concepts, Techniques and Tools*. Princeton university press.

Patton, A. J. (2006). Estimation of multivariate models for time series of possibly different lengths. *Journal of applied econometrics*, 21(2):147–173.

Sklar, M. (1959). *Fonctions de répartition à n dimensions et leurs marges*. Université Paris 8.

Tse, Y. K. (2000). A test for constant correlations in a multivariate garch model. *Journal of econometrics*, 98(1):107–127.

Xie, L., Kennedy, L., Chang, S.-F., Divakaran, A., Sun, H., and Lin, C.-Y. (2005). Layered dynamic mixture model for pattern discovery in asynchronous multi-modal streams [video applications]. In *Acoustics, Speech, and Signal Processing, 2005. Proceedings.(ICASSP'05). IEEE International Conference on*, volume 2, pages ii–1053. IEEE.

Yu, J. (2012). A particle filter driven dynamic gaussian mixture model approach for complex process monitoring and fault diagnosis. *Journal of Process Control*, 22(4):778–788.

Zucchini, W. and MacDonald, I. L. (2009). *Hidden Markov models for time series: an introduction using R*. CRC press.